AD-A138 533 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 7 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 01 OCT 83 CCTC-CSM-UM-246-82-CHG-7 1/3

UNCLASSIFIED F/G 9/2 NL

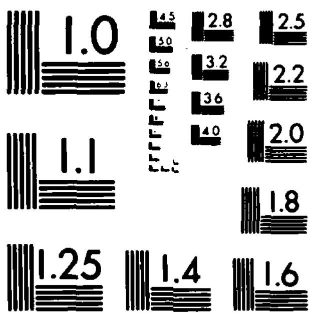

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

DEFENSE COMMUNICATIONS AGENCY
COMMAND AND CONTROL TECHNICAL CENTER
WASHINGTON, D. C. 20301

(2)

IN REPLY REFER TO: C751

1 October 1983

TO: RECIPIENTS

AD-A116898

SUBJECT: Change 7 to Computer System Manual CSM UM 246-82, Generalized Monitoring Facility

1. This is Change 7 to Computer System Manual CSM UM 246-82, Generalized Monitoring Facility, dated 1 May 1982. Remove obsolete pages and destroy them in accordance with applicable security regulations and insert new pages as indicated below:

Remove Pages	Insert New Pages
v	v
xi	xi
xii	xii
xviii	xviii
xix	xix
	xix.1
	xix.2
xxiii	xxiii
xxiv	xxiv
xxvi	xxvi
	xxvi.1
2-11 through 2-20	2-11 through 2-20
	2-21 and 2-22
4-1 through 4-4	4-1 through 4-4
5-25 and 5-26	5-25 and 5-26
5-27.4 and 5-27.5	5-27.4 and 5-27.5
5-27.14 through 5-27.19	5-27.14 through 5-27.19
5-41 and 5-42	5-41 and 5-41.1
	5-41.2 and 5-42
6-5 through 6-10	6-5 through 6-10
6-15 through 6-18	6-15 through 6-18
6-21 through 6-24	6-21 through 6-24
6-24.1 and 6-24.2	
6-25 through 6-36	6-25 through 6-36
6-41 and 6-42	6-41 and 6-42
6-47 and 6-48	6-47 and 6-48
6-53 and 6-53.1	6-53 and 6-53.1
7-3 through 7-6	7-3 through 7-6
7-7 and 7-7.1	7-7 and 7-7.1
7-7.2 and 7-8	7-7.2 and 7-8
7-9 and 7-10	7-9 and 7-10

7-15 and 7-16	7-15 and 7-16
7-17 and 7-18	7-17 and 7-17.1
	7-17.2 and 7-17.3
	7-17.4 and 7-18
7-19 and 7-20	7-19 and 7-20
7-21 and 7-22	7-21 and 7-21.1
	7-21.2 and 7-22
7-23 through 26	7-23 through 7-26
7-31 and 7-32	7-31 and 7-32
7-39 and 7-40	7-39 and 7-39.1
	7-39.2 and 7-40
7-47 and 7-48	7-47 and 7-48
7-48.1 and 7-48.2	7-48.1 and 7-48.2
7-49 through 7-52	7-49 through 7-52
7-55 through 7-58	7-55 through 7-58
8-5 and 8-6	8-5 and 8-6
	8-6.1 and 8-6.2
8-7 and 8-8	8-7 and 8-8
8-9 and 8-10	8-9 and 8-9.1
	8-9.2 and 8-10
8-11 and 8-12	8-11 and 8-12
8-13 and 8-14	8-13 and 8-13.1
	8-13.2 and 8-14
8-15 through 8-20	8-15 through 8-20
	8-20.1 and 8-20.2
8-21 through 8-24	8-21 through 8-24
8-25 and 8-26	8-25 and 8-25.1
	8-25.2 and 8-26
8-27 and 8-28	8-27 and 8-28
	8-28.1 and 8-28.2
8-29 through 8-36	8-29 through 8-36
8-51 and 8-52	8-51 and 8-52
11-7 and 11-8	11-7 and 11-8
12-41 and 12-42	12-41 and 12-42
	12-43 and 12-44
Complete Section 14	Complete Section 14

2. The effective date of these change pages is 1 October 1983.

3. When this change has been made, post an entry in the Record of Changes and file this letter before the title page.

FOR THE DIRECTOR:

Samuel D. Pucciarelli

144 Enclosures

SAMUEL D. PUCCIARELLI
Assistant Deputy Director
for Computer Services

Section		Page
7.9	Tape Error Aborts	7-58
8.	CHANNEL MONITOR DATA REDUCTION PROGRAM (CMDRP)	8-1
8.1	Introduction	8-1
8.2	Data Collection Methodology	8-1
8.3	Analytical Methodology	8-3
8.4	Data Reduction Methodology	8-7
8.5	CMDRP Output	8-8
8.5.1	System Configuration and Channel Usage Report (File 57)	8-8
8.5.2	System Summary Report (File 57)	8-10
8.5.3	System Traces Captured by Monitor Report (File 57)	8-13
8.5.4	Channel Status Changes Report (File 57)	8-13
8.5.5	Physical Device, Device ID Correlation Table (File 57)	8-13
8.5.6	Channel Statistics Report (File 57)	8-13
8.5.6.1	Channel Busy and Device Busy Report	8-13
8.5.6.2	Channel Busy and Device Free Report	8-13
8.5.6.3	Channel Free and Device Busy Report	8-13.1
8.5.6.4	Channel Free and Device Free Report	8-20
8.5.6.5	GEPR Connect Report	8-20
8.5.6.6	Lost Interrupt Report	8-20
8.5.6.7	Device ID STIOS Not Connected Report	8-20
8.5.6.8	Entries Still in Queue Report	8-20.1
8.5.6.9	Device Free But Has a Queue Report	8-20.1
8.5.6.10	500 Disk Drive Report	8-20.1
8.5.7	Idle Monitor Report (File 57)	8-20.1
8.5.8	Proportionate Device Utilization Report (File 57)	8-27
8.5.9	Queue Length and Queue Time Histograms (File 57)	8-29
8.5.10	Service Time Histograms (File 57)	8-34
8.5.11	Activity Statistic Report (Files 23 and 24) . . .	8-37
8.5.12	Job Conflict Report (Files 31, 32, 33, 34)	8-37
8.5.13	Special Job Processing Report by Device (File 32)	8-41
8.5.14	Special Job Processing Report Per 10 Minutes (File 32)	8-41
8.5.15	Special Processing Messages	8-44
8.6	Default Option Alteration	8-45
8.6.1	Job Device Conflict Report (Action Code QDEV)	8-46
8.6.2	Program Debug Options	8-46
8.6.2.1	Program Number Debug (Action Code DPRG)	8-46
8.6.2.2	Device Debug (Action Code DDEV)	8-46
8.6.2.3	Queue Location Debug (Action Code DQUE)	8-46
8.6.2.4	Random Access File Debug (Action Code DEBUG)	8-46

Section		Page

Section		Page
14.6.7.1	Background	14-49
14.6.7.2	Report Organization	14-51
14.6.7.3	TSS Tunable Parameters	14-51
14.6.7.3.1	Priority 'B' Processing Parameters	14-52
14.6.7.3.2	Placement of TSS Files	14-53
14.6.7.3.3	TSS Swap File Processing	14-55
14.6.7.3.4	Subsystem Accounting	14-57
14.6.7.3.5	UST Memory Management	14-57
14.6.7.3.6	Subsystem Memory Management	14-60
14.6.7.3.6.1	Subsystem Memory Allocation Processes	14-60
14.6.7.3.6.2	Settable Parameters	14-61
14.6.7.3.7	TSS Executive Processing	14-66
14.6.7.3.8	Subdispatch Queue Processing	14-68
14.6.7.4	Identification of TSS Response Time Problems	14-70
14.6.7.4.1	General	14-70
14.6.7.4.2	Analysis of TEARS RESPONSE Phase Reports	14-71
14.6.7.4.2.1	Timesharing Reduction Event Log	14-71
14.6.7.4.2.2	Response Times for all Users	14-73
14.6.7.4.2.3	Response Times For User Not Requesting More Core	14-73
14.6.7.4.2.4	Response Times For Users With Core Request During Line Idle	14-78
14.6.7.4.2.5	Total Time in Subdispatch Queue Report	14-78
14.6.7.4.2.6	Time in Subdispatch Queue Waiting Service	14-83
14.6.7.4.2.7	Processor Time In Subdispatch	14-83
14.6.7.4.2.8	CPU Monitor Driven Reporting	14-83
14.6.7.4.2.9	TSS Subtraces Encountered Report	14-83
14.6.7.4.2.10	Derails Report	14-90
14.6.7.4.3	Analysis of TEARS Emulation Phase Reports	14-90
14.6.7.4.3.1	Exception Message Report	14-93
14.6.7.4.3.2	USERS Map Report	14-96
14.6.7.4.3.3	TSS Core Map Report	14-96
14.6.7.4.3.4	ERROR MESSAGE Report	14-96
14.6.7.4.3.5	INTERTRACE Reports	14-96
14.6.7.4.3.6	Memory Activity Report	14-96
14.6.7.4.3.7	SUBDISPATCH Reports	14-97
14.6.7.4.3.8	Users Swaps (Swap Rate) Report	14-97
14.6.7.4.3.9	User Swaps (Swap Amount)	14-97
14.6.7.4.3.10	User Swaps (Duration) Report	14-97
14.6.7.4.3.11	User I/Os - Duration Report	14-97
14.6.7.4.3.12	In SMC Wait (Duration)	14-97
14.6.7.4.4	Resolution of Response Time Problems	14-97
14.6.7.5	Analysis of Response Time Problems	14-102
14.6.7.5.1	Priority 'B' Processing Problems	14-103
14.6.7.5.2	TSS Executive File Problems	14-103

Section		Page

Figure		Page
6-3	Standard Plot	6-13
6-4	PLOT Action Code Format	6-14
6-5	TIME Action Code Format	6-16
6-6	ALLOFF/ALLON Action Code Format	6-17
6-7	System Bottleneck Chart	6-20
6-8	JCL to RUN MUDRP	6-22
6-9	MUM Title Page Report - Idle Monitor Active	6-25
6-10	MUM Title Page Report - Idle Monitor Off	6-27
6-11	System Program Load	6-30
6-12	Standard Histogram Report	6-32
6-13	Out-of-Range Histogram	6-33
6-14	THIS FIGURE DELETED	
6-15	Activity Resource Usage Report	6-42
6-16	SNUMB IDENT Report	6-44
6-17	Memory Map Report	6-45
6-18	Demand List Report	6-49
6-19	Abort Report	6-51
6-20	Jobs Out of Core Report	6-52
6-21	Excessive Resource Usage Report	6-54
6-22	Allocation Status Report	6-55
6-23	Standard Plot	6-57
6-24	Memory Statistics Report	6-59
6-25	Special Job Memory Demand Report	6-60
6-26	Special Job Memory Size Report	6-61
6-27	Distribution of Urgency Over Time Report	6-61.1
6-28	Zero Urgency Job Report	6-61.2
7-1	System Configuration and Channel Usage Report	7-6
7-2	MSM System Summary Report	7-8
7-3	System Traces Captured by Monitor Report	7-11
7-4	Channel Status Changes Report	7-12
7-5	Physical Device, Device ID Correlation Table	7-13
7-6	Device Space Utilization Report	7-14
7-7	Device Seek Movement Report	7-16
7-7.1	Space Utilization Report for Device 9	7-17.1
7-7.2	Space Utilization Report for Device 10	7-17.2
7-7.3	See Movement Report for Device 10	7-17.3
7-8	Head Movement Efficiency Report	7-18
7-9	System File Use Summary Report	7-20
7-10	Individual Module Activity Report	7-22
7-11	SSA Module Usage Report by Job	7-24
7-12	File Code Summary Report	7-25
7-13	Cat/File String Report	7-27
7-14	Connect Summary Report by USERID/SNUMB	7-30
7-15	Activity Summary Report	7-31
7-16	Device Area File Code Reference Report	7-33
7-17	Device File Use Summary Report	7-34
7-18	Chronological Device Utilization Report	7-35
7-19	FMS Cache Report	7-37
7-20	Proportionate Device Utilization Report	7-38

Figure		Page

Table 2-1. GMC Release Dependent Parameters
(Part 1 of 11)

Element	Program	LINE #	Variable	Explanation
1	GMF.TOP	100	SYS64	Used to control conditional assembly of GMC set=1 for W6.4(2H) release set=0 for W7.2(4J) release
2		10240-10740	-	Code in this area searches for trace processing within the dispatcher. Trace code must be within 500 octal locations of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACE within the dispatcher, which is where the trace processing code is located. The code being searched for is an LDAQ;STAQ;TRA 0,1. If these instructions are not found, GMF will abort with an NF abort. In order for GMF to operate, the $TRACE card in the boot deck must not *specify no tracing*. At a minimum, the $TRACE card must request at least a single trace to be captured. It is recommended that if a site wants to minimize trace logic overhead, but at the same time be able to run GMF without altering the boot deck, that the following trace card be used: $ TRACE 777777777777,643777777777
3		10760-11180	-	Code in this area is used to make a correction to accounting processing, if the correction has not already been made via patches. The code is searched for within 500 octal locations of .MIOS entry point. The code searched for is SBLA TRREG+7,$;ARL 12; ADLA .CRTOD,7. The ARL is changed to an ARS.
4	*MUM.INIT*	500-740	-	Code in this area searches the dispatcher for generation of trace types 10 and 11. The code searches

Table 2-1. (Part 2 of 11)

Element	Program	LINE #	Variable	Explanation
				for an ORA =010,DL instruction and an XED .CRTRC instruction 3 locations in front of the ORA instruction. If the XED is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the XED or NOP is found, the GMF will abort with a Z0 abort. The same logic is repeated, except this time an ORA =011,DL is searched for. These instructions must be located within 1500 octal locations of entry point 15 decimal of .MDISP. This entry point should be the address of location TRACE within the dispatcher.
5	MSM.INIT CM.INIT CAM.INIT	400-600 380-570 220-420	-	Code in these areas searches .MFALT for generation of trace type 15. The code searches for an ADLA =015,DL instruction and an XEC .CRTRC+2 instruction 6 locations in front of the ADLA instruction. If the XEC instruction is not found, then a NOP instruction is looked for and will be replaced with an XEC instruction. If neither the NOP or XEC is found, the GMF will abort with a Z4 abort. These instructions must be located within 1000 octal words of the entry point of .MFALT.
6	MSM.INIT CM.INIT	610-820 580-790	-	Code in these areas searches .MIOS for generation of trace type 7. The code searches for an ORA 7,DL instruction and an XED .CRTRC instruction 7 locations in front of the ORA instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z1 abort. These instructions

Element	Program	LINE #	Variable	Explanation
				must be located within 7000 octal locations of the entry point of .MIOS.
7	CM.INIT	1690-1890	-	Code in this area searches .MIOS for generation of trace type 4. The code searches for an ARL 18 instruction and an XED instruction 3 locations in front of the ARL instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z3 abort. These instructions must be located within 1000 octal locations of the entry point of .MIOS.
8	CM.INIT	1900-2100	-	Code in this area searches .MIOS for generation of trace type 22. The code searches for an ARL 12 instruction and an XED instruction 3 locations in front of the ARL instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z2 abort. These instructions must be located within 6000 octal locations of the entry point of .MIOS.
9	CAM.INIT	1110-1520	-	Code in this area searches .MDNET to locate the CMBXO table. The code searches for a BCI 1,C instruction within 5000 decimal locations of the entry point of .MDNET. If GMF cannot locate this instruction, GMF will abort with an NS abort. If this instruction is found, then a BCI 1,TSS instruction is searched for within 30 decimal locations of

Table 2-1. (Part 4 of 11)

Element	Program	LINE #	Variable	Explanation
				the previous instruction. If the second instruction is not found, GMF will abort with an NT abort. If this instruction is found, then an LDA T.55,2 instruction (000001235212) is searched for within 10 decimal locations of the previous instruction. If this instruction is not found, GMF will abort with an NU abort.
10	IDL.INIT	300-1620	-	Code in this area searches for the generation of a series of trace types using the same logical construct as previously described for the other INIT sections. Following is a list of instructions searched for: (1) LDX2 .CRTRC+3 and an XEC .CRTRC+2 3 instructions in front (2) LDX2 .CRTRC+3 and an XEC .CRTRC+2 3 instructions in front (3) ARL 12 and an XED .CRTRC 3 instructions in front (4) ORA =016,DL and an XEC .CRTRC+2 3 instructions in front (5) ORA =037,DL and an XED .CRTRC 4 instructions in front (6) ORQ =013,DL and an XEC .CRTRC 7 instructions in front (7) ORA =021,DL and an XEC .CRTRC+2 4 instructions in front These areas of code search for generation of trace types 0,3,22,16,37,13 and 21, respectively, and if appropriate code is not found, the following respective aborts will occur: Z6,Z6,Z7,Z8,Z9,ZA and ZB. The modules which are searched are listed below: (1) 1000 octal locations from the entry point of .MIOS (2) 1000 octal locations from the

Table 2-1. (Part 5 of 11)

Element	Program	LINE #	Variable	Explanation
				entry point of .MIOS (3) 6000 octal locations from the entry point of .MIOS (4) 1000 octal locations from the entry point of .MFALT (5)-(7) 6000 octal locations from the entry point of .MDISP
11	CPU.PAT	300-440	-	Code in this area searches for an ASA .SALT,5 instruction in the dispatcher. Code must be within 300 octal locations after the address specified by entry point 20 decimal of the dispatcher. This entry point should contain the address of location DWAIT within the dispatcher. The trace code to be inserted by GMF indicates that a job is being taken out of processing. If the ASA instruction is not found, an N1 abort will occur.
12		460-590	-	In a WW7.2 (4JS) system, code in this area searches for an STQ .QTOD,4 instruction. Search area is the same as described for line 300. This trace indicates that subdispatching has finished using the processor. If the STQ instruction is not found, an N8 abort will occur.
13		640-800	-	If the dispatcher queue option of the CPU Monitor is activated, code in this area searches for an ORSA .STATE,4 instruction followed by an LDA .STATE,4 instruction. Code must be within 100 octal locations after the address specified by entry point 7 of the dispatcher. This entry point should contain the address of location DSPQH within the dispatcher. This trace indicates that a job is being placed into the

Element	Program	LINE #	Variable	Explanation
				dispatcher queue. If these instructions are not found, an NA abort will occur.
14		810-970	-	If the dispatcher queue option of the CPU Monitor is activated, code in this area searches for a LCQ=010001,DL instruction followed by an ANQ .STATE,5 instruction. These instructions must be within 1500 octal locations after the address specified by entry point 1 of the dispatcher. This entry point should contain the address of location DSP within the dispatcher. This trace indicates that a job is being taken out of the dispatcher queue and placed into execution. If these instructions are not found, an NB abort will occur.
15		970-1310	-	In order to implant its special hooks, the CPU Monitor must modify the dispatcher and, therefore, it requires eight words of patch space. If the dispatcher queue option of the CPU Monitor is activated, then 16 words, instead of the normal 8, are required. This patch area must be within 200 octal locations in front of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACF within the dispatcher.
16		2190-2320	-	If sufficient patch space was not available in the standard patch area, the CPU Monitor will attempt to locate patch space in a specially defined user patch area. This search will take place only if bit 2 of word 0 within the dispatcher is set. This patch space should be within 200 octal locations after the

Element	Program	LINE #	Variable	Explanation
				address ILIST within the dispatcher. The address of ILIST is found in word 7 of the dispatcher. If sufficent patch space is not found, an N2 abort will occur.
17		1950-2360	-	Code in this area searches .MFALT for GATE LOOP timing logic and attempts to correct logic errors that exist within the code. The code searches for an ARL 12 and ASA.CROVH,7 instruction within 2000 octal locations of entry point 13 of .MFALT. This entry point should be the address of the location BOOT within .MFALT. Twelve instructions in front of the ARL instruction, a LDX6.CRPRG,7 instruction is searched for. If this instruction is found, the code will next obtain the address stored in the upper half of the word located two locations in front of the entry point of .MFALT. This should be the address of the MMECNT table. The code will add 46 decimal to this address and should now contain the address of a patch table. The code will search for nine free words, and if found will correct the logic error. If patch space is not found, the code will remain unaltered.
18 19	CAM.PAT	170-290	-	Code in this area searches for an LDQ M.LID,3 instruction, followed by an ANQ=0077777,DU instruction in module DNWW/DNET. The search for this code begins at the address specified by entry point -8 of DNWW/DNET and continues until the code is located. This entry point should contain the address of location NRQWT-DNET within DNWW/DNET. If the instructions are not found, an N3 abort will occur.

Table 2-1. (Part 8 of 11)

Element	Program	LINE #	Variable	Explanation
20		350-500 and 690-820	-	In order to implant its special hooks, the CAM must find 8 words of patch space. Even though the CAM is modifying module DNET, it will use the patch space available in the .MDISP module. This search for space is identical to that for CPU.PAT at lines 970-1310 and lines 2190-2320. If patch space is not found, an N4 abort will occur.
21	MSM.PAT	200-240	-	Code in this area searches the dispatcher for SSA cache code. If bit 4 of word 0 in the dispatcher is set, then cache is available. If this bit is not set, then no further searches are performed.
22		250-380	-	Code in this area searches the dispatcher for the location of DBASE. The address at entry point -2 is obtained. This address points to the location ILIST in the dispatcher. At location ILIST, a series of addresses are stored and MSM.PAT searches this list for the address DBASE. If DBASE is not found, an N5 abort will occur.
23		410-620	-	Code in this area searches for an AOS .CRTDL and an AOS .CRTBH instruction. This code needs to be within 300 octal locations after the address DBASE within the dispatcher. If this code is not found, an N5 or N6 abort will occur.
24		670-790 and 1120-1250	-	In order to implant its special hooks, the MSM Monitor must modify the dispatcher and, therefore, it requires 8 words of patch space. This search is identical to that for CPU.PAT at line 680 and lines 970-1310 and 2190-2310. If patch space is not found, an N7 abort will occur.

Table 2-1. (Part 9 of 11)

Element	Program	LINE #	Variable	Explanation
25	GMF.MON	1040	FMS1	Offset from entry point of .MFSIO which points to the word giving the absolute address of FMS catalog cache buffer. Used only in W7.2. Set to -13 decimal.
26		1050	FMS2	Offset from entry point of .MFSIO pointing to the word which gives the option selection for FMS catalog cache. Used only in W7.2. Set to -15 decimal.
27	MUM.T10	220	SYS64	See GMF.TOP
28		920	FIFO	Address of the FIFO buffer within PALC. It is used to search the JCT table of PALC. This includes adding in a 110 octal offset for the loading of PALC in W6.4. There is no PALC offset in W7.2.
29		5410	XPQ24	Location in CALC of the memory demand table. Set to octal 111.
30		5420	SLVSNB	Offset in slave prefix area of job SNUMB. Set to octal 36.
31		5430	MEMUSE	Offset in slave prefix area of loader memory use word. Set to octal 37.
32		5440	IDENT	Offset in slave prefix area of job IDENT. Set to octal 66.
33	CM.T07A	190	IDENT	Offset in slave prefix area of job ident. Set to octal 66.
34		210	SYS64	See GMF.TOP
35		9960-10130	-	Code in this area searches .MFSIO in order to gather statistics for FMS catalog cache analysis. All the following references are offsets from the entry point of .MFSIO:

Element	Program	LINE #	Variable	Explanation
				-12 # of cache hits -11 # of writes -10 # of reads 841 # of reads not in CC 842 # of 320-word reads 843 # of skips 844 # of cache clears 848 # of no hits 849 # of hits
36		10180-10250	-	Code in this area searches .MAS04 in order to gather statistics concerning the available space table utilization. All the following references are offsets from the entry point of .MAS04: -6 # of times buffer allocation attempted -5 # of times buffer busy -4 # of times AST was in memory -3 # of times AST in memory but busy
37		11440	FFCCC	Address in PALC where the file code is stored during GEFSYE processing. In PALC, variable is called FFC. Set to 6177 octal in W6.4, 13143 octal in W7.2.2, and 13201 octal in W7.2.3. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.
38		11450	SNUMBP	Address in PALC where the SNUMB is stored during GEFSYE processing. In PALC, variable is called SNUMB. Set to 35012 octal in W6.4, 2632 octal in W7.2.2, and 2642 octal in W7.2.3. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.

Table 2-1. (Part 11 of 11)

Element	Program	LINE #	Variable	Explanation
39		11460	ACT	Address in PALC where the activity number is stored during GEPSYE processing. In PALC, variable is called SACT. Set to 33231 octal in W6.4, 1051 octal in W7.2.2, and 1061 octal in W7.2.3. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in W7.2.

THIS PAGE LEFT INTENTIONALLY BLANK

SECTION 4. RESOURCE MONITOR DATA REDUCTION

The Resource Monitor Data Reduction (RMDR) is composed of three discrete programs. The first program (PSUMR) processes the SCF records producing an intermediate record. The second program is a daily monitor (DEMON) which maintains an on-line history file. The third program (RPTSUM) produces reports and prints the required plots. Figure 4-1 is an overview of the RMDR.

4.1 PSUMR

PSUMR selects the collector data records from an SCF input file. The records are sorted by system ID, date, time, and record subtype. The records are then processed in pairs to produce an intermediate record.

4.1.1 PSUMR Inputs. PSUMR requires one input file containing SCF records. Two optional inputs are the intermediate record - old master, and a parameter file.

4.1.1.1 SCF Records (File IN). Contains resource monitor collector (RMC) SCF records.

4.1.1.2 Intermediate Record - Old Master (File OM). This optional input will be copied to the output file before new records are added from the current jobs processing.

4.1.1.3 Parameter File (File PF). PSUMR will process the following parameter language statements:

CONFIG
SUMMARY
SNUMB

The format and use of the parameter language is found in section 4.4.

4.1.2 PSUMR Output. Three outputs are available from PSUMR: (1) listing of the parameter file if present; (2) intermediate record - new master file; (3) intermediate record - current file.

4.1.2.1 Parameter File Listing (File RP). This is a listing of the contents of file PF. Any of the statements processed which were invalid will be flagged. Invalid contents will not halt the job execution, but may affect how the data is summarized in the intermediate record.

4.1.2.2 Intermediate Record New - Master File (File NM). This file contains the contents of the old master file, if present, plus the records added from the current execution of the program.

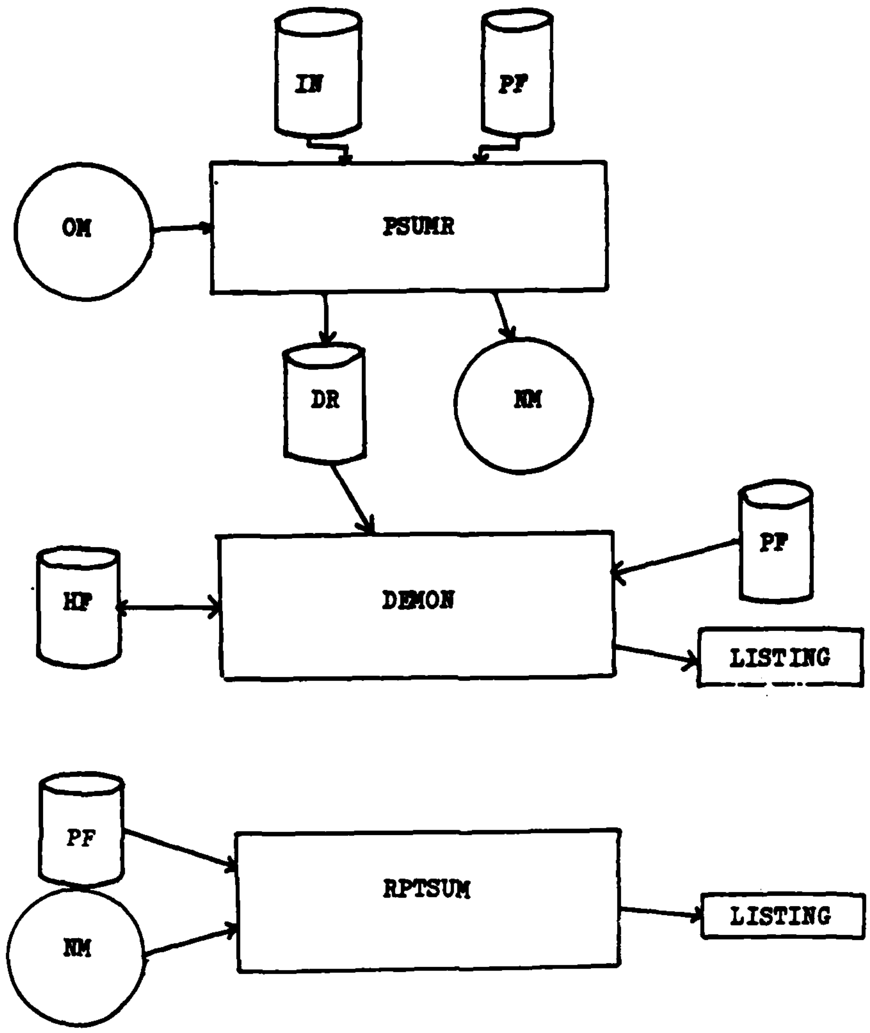

IN = SCF accounting records
PF = optional parameter file
OM = old master intermediate record
NM = new master intermediate record
DR = optional intermediate output for DEMON
HF = daily monitor history file

Figure 4-1. RMDR Overview

4.1.2.3 Intermediate Record - Current File (File DR). This file contains the intermediate records created during the current execution of the program. This file is used to pass the intermediate records to the daily monitor (DEMON), when present, in a subsequent activity.

4.1.3 PSUMR Deck Setup. The following control cards are required to execute PSUMR:

$	IDENT	ACCOUNTING INFORMATION	
$	USERID	USERID$PASSWORD/SCC	
$	PROGRAM	PSUMR,Dump	
$	LIMITS	20,20K,,5K	
$	PRMFL	**,R,R,B29IDPXO/NEWRMON/DDYN	
$	PRMFL	H*,R,R,B29IDPXO/NEWRMON/DDYN	
$	DATA	PF	optional parameters
$	TAPE	OM	optional old master intermediate record
$	TAPE	NM	new master intermediate record
$	FILE	DR	current file output
$	SYSOUT	RP	parameter file listing
$	TAPE	IN	SCF input file
$	FILE	S1,S1R,50R	sort files
$	FILE	S2,S2R,50R	
$	FILE	S3,S3R,50R	
$	FILE	S4,S4R,50R	
$	ENDJOB		

The sort files' size should be increased or decreased depending upon the input volume. Tape files may be replaced by permanent or temporary disk files. File DR may be null if the daily monitor does not follow in a subsequent activity.

4.2 DEMON

The daily monitor accepts the intermediate records created by PSUMR. A history file will be initialized (if current null) or updated from the data in the intermediate record. A matrix structure is built to summarize and store the data. The matrix contains 96 entries corresponding to the time period to be summed together. The entry size is variable depending upon the graphs to be produced. For DEMON, the default is all graphs. The history file contains a copy of this matrix structure. For an update run, DEMON copies the history file into core. A second matrix is built for the current day's data. Both matrices are then updated from the input. The history matrix is written to disk. Finally, the current input matrix is compared to the history matrix. If 25% of all data items are not within 25% of the history matrix value, the program generates a complete set of graphs to indicate a significant change from the history file value.

4.2.1 DEMON Inputs. DEMON has one required input and two optional inputs.

4.2.1.1 History File (File HF). The history file is required for DEMON execution. The history file is a random file which contains an image of the in-core matrix used for summarizing and comparing data.

4.2.1.2 Parameter File (File PF). The parameter file is optional. However, if PSUMR was supplied with a SUMMARY statement, the DEMON should be supplied the same summary statement for desired results. DEMON processes only a summary statement. The format and use of the parameter language statements is found in section 4.4.

4.2.1.3 Intermediate Records (File IN). This is an optional input. If present, the intermediate record data will be summarized, as described above, according to parameter setting marked on the intermediate record. The history file will be initialized or updated as it is appropriate. If file IN is not present, the history file data will be used to generate a set of graphs. This feature allows the user to get a current historical "picture" at any time.

4.2.2 DEMON Outputs. DEMON outputs are the updated history file and a listing file.

4.2.2.1 Updated History File (File HF). Refer to section 4.2.1.3.

4.2.2.2 Listing File (File RP). The output listing will contain a listing of the parameter file, if any, and any generated graphs. The graphs produced are listed and described in section 4.5..

4.2.3 DEMON Deck Setup. The following control cards are required to execute DEMON.

```
$     IDENT        ACCOUNTING INFORMATION
$     USERID       USERID$PASSWORD
$     PROGRAM      DEMON,DUMP
$     LIMITS       10,24K,,10K
$     PRMFL        **,R,R,B29IDPXO/NEWRMON/DDYN
$     DATA         PF        optional parameter file
$     TAPE         IN        optional intermediate file
$     PRMFL        HF,W,R, (history cat/file)
$     SYSOUT       RP
$     ENDJOB
```

File IN may be a permanent or temporary disk file. File HF is 24 random LLINKS.

4.3 RPTSUM

PSUMR has the ability to maintain a master file of intermediate records. RPTSUM will process this file to produce any number of reports. Each report

5.2.8 Idle Monitor. The Idle Monitor (IDLEM) is used to collect data concerning CPU activity. This monitor can only be used in conjunction with the MUM or the CM. It should not be activated if one of the two aforementioned monitors is not active. If the Idle Monitor is present on the R* file and active and if, in addition, the MUM or CM is not active, then the IDLEM will automatically be turned off. The user should read subsections 5.2.1 and 5.2.5 for information concerning the use of the IDLEM. A separate discussion of the format of the collected data records is contained in subsection 5.4.9.

This monitor generates an excessive number of traces and significant overhead. It should be activated only after the user ensures that the reports produced because of its presence are an absolute necessity for the evaluation. In most cases, this monitor should not be required.

5.2.9 Transaction Processing System Monitor. The GMC Transaction Processing System Monitor (TPSM) is used to collect data on the performance of the Navy Transaction Processing Executive (NTPE) System. A separate discussion of the format of the TPSM collected data records is contained in subsection 5.4.10. The reports containing data collected by TPSM are described in section 12.

When TPSM is active, the required traces must be enabled in the computer system boot deck on the $ TRACE card (see table 5-1). This is the only GMF monitor for which the minimum $TRACE card format described in subsection 5.1 is not sufficient for proper execution. A sample of the reports and run time procedures for the data reduction program can be found in the Transaction Processing System (section 12). The TPSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 9, 11, 19, 23, and 36. The complete process for generating an R* file is described in subsection 5.6.

NOTE: The TPSM cannot be run concurrently with the TSSM and should only be used on a WWMCCS WW7.2 system. It should not be used on a commercial release.

5.2.9.1 TPS Trace Collection. The TPSM is unlike most other GMC monitors in that monitoring of the Transaction Processing System is controlled via the operator console. Prior to collecting data, the user must alter the NTPE (see subsection 5.2.9.2) and must also create a usable GMC R* file (see subsection 5.6). Once these actions are performed and a GMC execution is started, the user must still perform one additional action before data collection can begin. The TPSM is turned on or off by the console operator via the TP MESS command. The operator must request "TP MESS". When the console responds "TP MESS?", the message "TRACE ON" is entered to start trace generation, or "TRACE OFF" to suspend trace generation. This procedure can be repeated as often as desired. The TPSM and the TSSM are the only GMC monitors that can be turned on or off while the GMC is physically executing.

5.2.9.2 Modifying the Transaction Processing System. To use the TPSM, the user must alter the Navy Transaction Processing System. NTPE is delivered as a set of alters to the commercial version of TPE (these alters are in "*C" form for use by the SCED utility program). Embedded in these original alters are the modifications to place the GMC trace points in the executive modules (see figure 12-21). All sections of code for the GMC support are under control of conditional assembly parameters. To enable the GMC code, the changes shown in figure 12-22 must be merged with the local file or changes developed during the site customization of NTPE. The changes shown in this figure have the effect of both correcting and enabling the GMF/TPEM collector.

5.2.10 Timesharing Subsystem Monitor. The Timesharing Subsystem Monitor (TSSM) is used to measure TSS performance. Section 15 details those reports available from the data collected by this monitor. This monitor should be used only on a WWMCCS WW7.2 system. If desired to be run on a commercial release, care should be exercised to ensure that all alters are located correctly (see subsection 5.2.10.1).

The TSSM causes the trace to be taken from many points in TS1; the collector builds its records which are then passed to the ER for buffering. An example of the record format appears in subsection 5.4.11. TSSM requires the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 10a, 11, 19, 23 and 36a. Subsection 5.6 shows how to generate an R* file.

9 TSSH RSPRM+6 LDA -1,DU Enter build mode

Data: UST address, program stack tally
Use: Cause entry to be made in stack of subsystem names; remove entry from stack when trace 13 is received

10 TSSH LODPRM+9 TNC SCERR6 EXEC primitive

Data: UST address, program stack tally, subsystem name (ASCII)
Use: Set flag to indicate actual code executed (as opposed to subsystems CARD, FORT, MDQ, etc.)

11 TSSH SYSPRM+4 LDA .LBUF,2 SYSTM primitive

Data: UST address
Use: Clear stack of subsystem names

12 TSSI CU1CF+17 EAX1 .LCALS,2 Log-off

Data: UST address
Use: Terminate processing of session, release UST address

13 TSSI STARTP+3 LDA 0,3 Command received

Data: UST address, command text (ASCII)
Use: Record last user build mode command for session snapshot reports; end build mode; clear subsystem name stack if command is "NONE" and does not come from DRL CALLSS or DRL T.GOTO

14 TSSJ LOGON+1 CANA =07777,DL Periodic check

Data: Line ID if new user waiting
Use: Produce report entry if time interval between trace 102 and this one too great

15 TSSJ LOG11+3 STQ BYTIME GEWAKE - no users

Data: Length of GEWAKE (clock pulses)
Use: Force clearing of tables in data reduction program; opens interval closed by trace 61

16 TSSJ LN100+2 STA .LKST,2 Break or disconnect

Data: UST address, GEROUT status (binary)
Use: Set flag for reconnect if status 2 (disconnect); simulate input complete if break and if I/O complete trace has been processed (user initiates request for service with break as in DJST)

17 TSSJ .SSDSP+7 TNC TYTSS GEWAKE until subdispatch done

Data: Length of GEWAKE (clock pulses)
Use: Formatted dump or accounting for GEWAKEs; opens interval closed by trace 61

18 TSSJ TYTSS STQ SLEEP GEWAKE with subdispatch busy

Data: Length of GEWAKE (clock pulses)
Use: Formatted dump or accounting for GEWAKEs; opens interval closed by trace 61

19 TSSJ QSPEC EAA SPEC All Points Bulletin or remote I/O courtesy call

Data: UST address, flags (.LFLG2) for CRUN/DRUN, GEROUT status
Use: Completes terminal I/O started by trace 100

20 TSSJ STGCC+5 CMPA 4,DL Build mode input received

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O started by trace 100

21 TSSJ DSTAT+4 LDA .LIOST,2 SY** I/O complete

Data: UST address
Use: Completes disk I/O started by trace 24

22 TSSJ RECON2+1 CANA .FK19,DL Place user in reconnect mode

Data: UST address, flag for data in transmission
Use: Positive assurance that session is being put in hold (not given by trace 16)

59 TSSJ PPTCC2+2 STZ .LCC,2 Disk I/O for tape mode complete

Data: UST address
Use: Completes disk I/O started by trace 24

104 TSSK RETSBS+1 SZN .LSFLG DRL processing complete

Data: UST address, A-register status if DRL T.SYOT or DRL TASK, Q-register status if DRL SPAWN or DRL PASFLR, DRL number if status reported
Use: End DRL processing state; report unusual status conditions

TSS passes flag to GMC to write initialization records or GMC writes initialization records after it has recovered from a lost data condition)
If TRACE OFF, flag (=0)
If bad console input, no trace

Use: Initialize tables in data reduction program

91	TSSM	RETS3X+6	CMPX1 1,DU	Make subdispatch entry

Data: UST address
Use: Open interval closed by trace 88

92	TSSN	GUST1+2	CMPA .LTSRI,DL	Process log-on request

Data: Line ID (octal 2020 if deferred), reject flags (.TLFLG bit 32 or 35)
Use: Open interval closed by trace 1

93	TSSN	GUST21+15	TRA GUSTR-1	Reject user - bad line status

Data: Line ID
Use: Report along with log-on rejections (this is an exotic line condition trace)

94	TSSN	GUST2A+4	CMPQ =030000,DU	Check for VIP as terminal type

Data: Terminal ID, type, number of VIPs allowed (.T760), number logged on (.TL760)
Use: Explain trace 96 if it occurs

95	TSSN	GUST4+4	CMPQ GUSTT	Check UST wait time

Data: Line ID, flag if wait greater than 16 seconds
Use: If flag set, report as log-on reject

96	TSSN	GUSM+5	STA 2,3	Reject user

Data: Line ID
Use: Report trace if not already done

97	TSSN	MUST3+2	AWDX ,2,0	UST compression

Data: Old, new UST addresses
Use: Update UST-oriented tables with new address

98	TSSN	ASGGC1+4	ASQ 1,0	UST area increase by 1K

Data: None
Use: Decrement amount of user memory available

99 TSSN RELW03+18 ASA 1,0 UST area decrease by 1K

Data: None
Use: Increment amount of user memory available

100 TSSN RIO.1A+2 LDA .LBUF,2 Terminal I/O request

Data: UST address, GEROUT opcode, CRUN flags (.LFLG2)
Use: Open interval closed by trace 19, 20, 84, 86, 87, or 103

101 TSSN R.SFOK+2 LXL1 -1,6 Process command file $*$ function

Data: UST address, text (ASCII)
Use: Provide information to explain changes in .LFLG2 or to explain trace 6

76 TSSN TRMCF+7 EAQ 0,AU Cancel CRUN mode

Data: UST address, flags (.LFLG2)
Use: Force clearing of subsystem name stack

102 TSSN IRINQ+4 LDX1 =0050000,DU Issue remote inquiry GEROUT

Data: None
Use: Open interval closed by next trace; if long time, indicates roadblock of TSS in .MROUT because a table is full

103 TSSO 760CC+3 CMPA 4,DL VIP input complete

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O started by trace 100

5.2.10.2 Formats of TSS Traces. The collection routine executing within TSS passes two-word traces to GMC through the dispatcher trace mechanism. In all but four trace subtypes, GMC stores the A and Q registers after a record control word and RSCR time to form a four-word logical record. For subtype 24, GMC edits the Q register before making a trace. For subtype 32, GMC additionally retrieves the USERID and appends it to form a 6-word logical record. Subtype 105 is generated internally by GMC when it senses a period of no TSS activity. Logical records generated as a result of receipt of a subtype 90 trace may have 3 different formats. The first format is written whenever TSS passes a subtype 90 trace to GMC indicating that the TSS traces have been turned off. GMC passes this trace to the data tape in the standard format described above. The following 2 formats apply to a group of logical records written by GMC whenever it recovers from a lost data condition or whenever TSS passes a subtype 90 trace indicating that TSS traces have been turned back on. In the latter case, GMC formats a series of logical records; one 16-word record is written for each user logged onto TSS to give such attributes as UST address, line ID, USERID, and current program stack. A four-word record which indicates the TSS swap area limits is the last record written by GMC in the sequence of records. If no users are logged on to TSS and if TSRI is not running, only the four-word record is written. The Monitor source code gives the formats of all TSS record subtypes.

5.2.10.3 Installation Procedures

5.2.10.3.1 *Description of Monitor Software.* The TSS Monitor program element executes as a TSS master subsystem. When the master user logs on and requests the program via "SYST GMF", it first performs address relocations needed to convert relative offsets into actual slave addresses. Then it verifies that the instructions at the trace points match those in the monitor coding (the originals have transfer instructions patched over them when the monitor is executing). Any mismatches found are reported on the master terminal and verification continues. If any mismatches are found, the subsystem terminates without making any modifications to TSS. If verification succeeds, the master subsystem copies part of itself (about 1K memory) into an area of module TSSO which was reserved by means of patches on the TSS INIT file. The origin of this area is the UST address origin which TSS would have used were the patch not applied. Next, the master subsystem applies Execute Double (XED) instructions to most trace points to save the return address and indicator register. In a few rare cases, unconditional transfer instructions are used. The master subsystem terminates at this point. As TSS continues execution, traces will be formatted, but a switch prevents GCOS traces from being written. If the console operator enters "TS1 TRACE ON", GCOS traces will be written from TSS, and an additional overhead will be imposed on TSS. The console operator can stop generation of TSS traces with "TS1 TRACE OFF". The traces can be turned on and off multiple times.

5.2.10.3.2 Software Installation. When TSS is loading the monitor subsystem, it must be able to access the program element using a MME GECALL to a BCD name of .TSGMF. The program element may reside either on a system file defined in the EDIT and FILES sections of the boot deck or on a permanent file (B29IDPXO/GMFCOL/TSS/TSSMON) dynamically accessed during TSS startup. Use of a permanent file requires additional patches in the TSS INIT file. This is the current method of implementation because it does not require changes to the GCOS startup deck. The job stream used to create the program element is located on file B29IDPXO/GMFCOL/TSS/TSGMF.

If the system file option is to be used, the $ PRMFL Q* card on file TSGMF must be replaced with a $ TAPE Q*,X2D,,,,TSSMON card. The startup file must be defined in the EDIT section of the boot deck as follows:

```
$      FILDEF   ST1,TSSMON,12/O,SYS,1T1
```

The tape drive name 1T1 must be appropriate to the hardware configuration. If this startup file is appended onto an existing edit tape, replace "1T1" with "*". In the FILES section, insert the following card in front of existing $ SYSTEM cards:

```
$      SYSTEM   TSSMON
```

If a permanent file is used, no changes need to be made to the boot deck, and no system interruption occurs during installation of the TSS Monitor.

5.2.10.3.3 Software Activation. The purpose of this section is to describe how TSS builds tables so that the master user can find a subsystem named "GMF".

5.2.10.3.3.1 Overview of TSS INIT File Changes. The TSS INIT file is a quick access permanent file named TS1 normally residing under USERID OPNSUTIL, the default USERID for TSS and patchable as symbolic location .TUSER. File TS1 is read as soon as TSS starts in order to pass parameters for the current loading of TSS or to allow symbolic specification of site option patches such as the maximum number of concurrent users. The TS1 file has two sections: $INFO to symbolically define site option parameters, and $PATCH to apply patches to TSS beyond those already in the PATCH section of the boot deck. Installation of the TSS monitor requires that patches be placed at the end of the TSS INIT file. These patches are located on file B29IDPXO/GMFCOL/TSS/TSS.PAT (updated for SRW7.2.3).

5.2.10.3.3.2 Definition of the Master Subsystem Name. The following three patch cards, included in file TSS.PAT, overlay an unused program descriptor in TSSA to define a subsystem named GMF with an edit name of .TSGMF and attributes .BPRIV (can execute privileged DRLs, not currently used), .BMAST (master subsystem), and .BMASX (permission to alter TSS executive with a DRL instruction, not currently used):

```
6735   OCTAL   147155146040   GMF                 .MTIMS
6736   OCTAL   336362274426   .TSGMF              .MTIMS
6737   OCTAL   40003          MASTER SUBSYSTEM    .MTIMS
```

5.2.10.3.3.3 Definition of the UST Origin. The following patch, included in file TSS.PAT, disable the check for overlaying VIP code in TSSO when no VIPs are configured. The patch NOP's out a conditional transfer instruction around an instruction which loads a UST origin corresponding to the case when VIPs are configured. The letter "O" in these patches means that the patch locations are with respect to the origin of TSSO, not to slave address 0 of TSS as in the subsystem attribute patches. The letter "R" before patch content indicates that the address field of the patch must be relocated to the beginning of the module (TSSO) identified in column 7 of the patch. The symbolic addresses for this patch is IN7630+2. The following patch completes the TSS INIT file when the TSS Monitor is loaded on a startup file:

```
5275  OOCTAL   11007           DON'T TRANSFER IF NO VIPS      .MTIMS
```

5.2.10.3.3.4 Installation from a Permanent File. The principle of this method is that if a job has a file code ** active, any MME GECALL will cause that file to be searched before the system files are searched. Patches on the TSS INIT file must access a permanent file before subsystem loading begins and release it after subsystem loading finishes. The TSS User Derail Loader (TUDL, SDN K79005) provides an example of how to access and release a system loadable file from within the TSS executive. The patches on the TSS INIT file for the TSS Monitor are compatible with TUDL because TUDL starts its work after the TSS Monitor has finished its work. TUDL may further relocate the UST origin upward, but TUDL uses the address stored by the TSS Monitor patches. A list of an entire $ PATCH section is given in the TSSM source code. The first four patches match ones described in earlier sections.

The bulk of the patches used to access a permanent file are placed in an area of TSSO which is reserved for UST space when the UST origin is not adjusted upward (actually, this space was needed in releases prior to W7.2.0 to prevent the generation of a UST for TSR1 from destroying the code at the end of TSS startup; with the TSS INIT file feature, enough code intervenes so that this buffer is not necessary and so that the UST origin can be moved up and not destroy the code). The first transfer into these patches is at offset 4677 in TSSO (symbolic offset DMYERR+2, instruction EAX5 1). Here, the USERID in the file structure defined in the patches is stored into the SSA of TSS at location .SUID. Then a MME GEMORE accesses the file. Use of .SUID makes FMS think that the file is being accessed by its owner and thus, no special permissions are needed. If the MME GEMORE is denied, a flag is set to 0 before control returns to the original TSSO coding from offset 2013 in the patches.

The second transfer, at octal offset 5272 (symbolic location IN7630+3) replaces the instruction defining the UST origin. The patch at octal offset 5271 remains intact. In the second group of patches, the permanent file must be released, the number of PATs decremented by one, and the USERID in .SUID set to zero. If the GEMORE in the first set of patches is denied, then the number of PATs is not decremented. This alteration of word .SNPAT in the SSA is necessary because releasing the file does not remove the file code from the SSA. The patch for defining the new UST origin is moved to octal offset 2031 in the patches, just before a transfer back to TSSO.

5.2.10.4 Production Use of the Monitor. The following steps must be taken to enable data capture:

1. Append patches on file B29IDPXO/GMFCOL/TSS/TSS.PAT to Timesharing INIT File (normally OPNSUTIL/TS1).

2. Run the file B29IDPXO/GMFCOL/TSS/TSGMF to create the new TSS subsystem file.

3. Start TS1.

4. Start a copy of GMC with the TSS Monitor active.

5. Log on to a master USERID and enter "SYST GMF" at the "*" prompt following log-on to install the hooks into TSS code. Traces will not start at this time. The master USERIDs assembled in TSS are MASA and MASB. Unless these are patched or redefined in the TSS INIT file, only a terminal ID designated as master in .MSECR may be used for this step.

6. Enter "TS1 TRACE ON" from the system console to start generation of traces from TSS.

7. Enter "TS1 TRACE OFF" from the system console to suspend generation of traces.

Steps (5), (6) and (7) execute independently of (4); however, use of step (6) without use of step (4) will cause unnecessary TSS overhead if traces are being generated and lost due to GMC not being in execution. Steps (6) and (7) may be repeated multiple times if traces are to be captured for specific periods of the day. (The companion data reduction program, described in section 15 cannot process GMC sessions longer than 9 hours).

5.2.10.5 Monitor Limitations. To obviate lockup fault, initial trace generation (at "TS1 TRACE ON", and following lost data) is inhibited until the number of TSS users falls below 50. Further, all trace generation is inhibited until the number of users falls below 100. (The companion data reduction program is limited, by parameter, to 50 active USTs).

4	0-29	CP time usage
	30-35	IOM command
5	0-5	Not used
	6-35	SNUMB
6	0-5	Device command Special note: A command of octal 72 to a permanent disk pack indicates that a pack exchange is in progress. The .MGP66 module issues another standby command to the device to which the permanent pack is to be moved. A special device name record should appear either in the current block on tape or at the beginning of the next block to confirm the pack exchange.
	6-17	DCW length
	18-35	File origin block number
7	0-13	File size
	14	Sysout flag
	15	Seek flag
	16-35	Seek address
8	0-5	Device command
	6-11	Device number
	12-17	IOM PUB number
	18-23	I/O Command
	24-26	Not used
	27-29	IOM number
	30-35	Record count
9	0-17	MBA of job issuing connect or zero if nonextended memory
	18-35	I/O queue address in SSA (absolute address if nonextended memory or if value less than 64K; relative address (to MBA) if extended memory and if value greater then 64K)
10	0-8	Not used
	9-17	Activity number
	18	Flag (=1 - I/O status is stopped)
	19-28	Not used
	29-34	I/O status (I/O queue word 0)
	35	Flag (=1 - system job)
11	0-35	.CRCOM (use only bits 18-29)
12	0-35	.CRCOM+2 (use only bits 18-29)
13	0-35	.SGCPT (use only bits 18-35)
14	0-35	.SRQCT (use only bits 0-17)
15	0-35	.SNIO (use only bits 18-35)
(words 11-15 are not generated if the IDENT for a job is reported)		
11-20		IDENT (words 11-14 described above are not collected)
21-22		USERID

5.4.3.2 MSM Special Record. During the execution of MSM a special record is written at preselected times during the monitoring session. These records are used to analyze SSA cache core, when configured. The format of this record is shown below. This record is based on data collected during the processing of GMC trace events 73 or 76.

THIS PAGE LEFT INTENTIONALLY BLANK

Word	Bits	Information
1	0-17	Size of record (=517)
	18-26	Not used
	27-35	Octal 7 (trace number)
2-516	0-35	a. Number of times each GCOS module 1-515 was in the SSA cache buffer b. Number of times each GCOS module 1-515 was loaded by an I/O because it was not in the SSA cache buffer.
517	0-35	Not used
518	0-35	Flag (=2 - case a. above) (=3 - case b. above)

5.4.3.3 Device Name Record. If either MSM or CM is active, the GMC writes a record which correlates device names to device addresses. The System Configuration Name table is processed sequentially as this record is formatted. Names for all disk devices are reported. In order to detect exchanges of mass storage devices, GMC periodically examines the device name table. If any changes have occurred, then another device name record is written. This record is variable in size, and recognized by the special format of the second word.

Word	Bits	Information
1	0-17	Size of record
	18-26	Not used
	27-35	Octal 7 (trace number)
2	0-35	Octal 535353535353
3-n	0	Flag (=1 - fixed device if mass storage) (set if bits 13 and 14 of the second word of the SCT entry for any mass storage device are both zero; in Shared Mass Storage environment, shared devices must be fixed)
	1-5	IOM number
	6-11	Channel number
	12-17	Device number
	18-35	Device name in BCD

5.4.3.4 FILSYS Catalog Structure Record. During the execution of the Mass Store Monitor, certain MME GEFSYEs GCOS Trace 15 data are collected concerning the catalog file string that is being referenced. The purpose of this data is to try and determine how many connects are being made because of the particular structure of a given catalog or file. This data is also used to provide the catalog file string name associated with the various user file codes that are reported by the Mass Store Monitor. MME GEFSYE traces are only processed if generated by a system program (program number less then 15 or FSTSLV). In addition, only the following GEFSYE's will be

Table 6-1. (Part 4 of 4)

ID Number/Name	Other Reports
37/PALC	Peripheral Allocator Report
38/ACTIVE	Activity Report/Excessive Resource Report/Abort Report/IDENT Report
39/MAP	Memory Map
47/OUT	Out of Core Report
---	Special Job Memory Reports
---	System Program Usage Report
---	Memory Statistics Report
---	Distribution of Urgency Over Time Report
---	Zero Urgency Job Report

cover a larger range of values. This change could be made via data cards and would not increase the size of the program.

The second method would involve increasing the size of the histogram by altering the value of TABSIZ. As long as the size requested does not exceed 50, this change can also be done via a data card. However, if an individual histogram needs to be larger than 50 buckets, the user will need to change the value of MXTBSZ. This change will require a change to source code, a recompile, and probably an increase in program size. All references to MXTBSZ must be altered. This would need to be done in the EDIT subsystem of Timesharing.

The remaining items that can be modified are the title and the vertical axis headers. The method for altering the histogram parameters is detailed in subsection 6.1.6. Table 6-2 shows the default values for all histograms.

6.1.4 Plot Options. There are three characteristics directly available to the user for each individual plot axis used.

The first characteristic, MAXNUM, is the maximum number of entries to be plotted on each vertical plot axis.

The second characteristic, YMAX, defines the upper limit of the horizontal display axis.

The third characteristic, YMIN, defines the lower limit of the horizontal display axis. The method for altering these values is explained in subsection 6.1.7. Table 6-3 shows the default values for all plots.

6.1.5 Default Option Alteration. The general format for an option request is as follows: the first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirements being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" typed on it. This card must be present whether or not any other input options are selected. Available actions with their (default) implications are shown in table 6-4. There is no order required for the options. In reading the following sections it should be remembered that the first card for any input option must be the action code specification with no other data present on the card.

Table 6-2. Default Values for Histograms (Part 1 of 2)

ID #	Low Value	Interval Size	Number of Buckets
1	4	4	50
2	0	50	50
3	0	250	50
4	0	1	50
5	4	4	50
6	0	1	50
7	0	5	50
8	0	200	50
9	0	200	50
10	.95	.1	50
11	4	4	50
12	0	10	50
13	0	1	50
14	0	1	50
15	0	1	50
16	0.0	5.0	50
17	0	10	50
18	200 *	2 *	50
19	0	20	50
20	0	1	50
21	0	1	50
22	0	8	50
23	0	25	50
24	0	25	50
25	50	2	50
29	0	1	50
30	0	1	50
31	0.0	5.0	50
32	0.0	5.0	50
33	0.0	5.0	50
34	0.0	5.0	50
35	5	5	50
36	5	5	50
40	0	1	50
41	0	1	50
42	0	10	50
43	0	1	50
44	0	1	50
45	0	10	50
46	0	1	50
48	0.0	5.0	50
49	0.0	5.0	50
50	0.0	5.0	50

Table 6-2. Default Values for Histograms (Part 2 of 2)

ID #	Low Value	Interval Size	Number of Buckets
51	0.0	5.0	50
52	0.0	5.0	50
53	0	1	50
54	0	250	50
55	0.0	5.0	50
56	0.0	5.0	50
57	0.0	5.0	50
58	0.0	5.0	50

* - The Low Value and Interval Size parameters for histogram ID #18 are determined by the program in a dynamic manner.

LOW VALUE = MEMORY CONFIGURED ON SYSTEM - PRECODED LOW VALUE (200)

INTERVAL SIZE = ([MEMORY CONFIGURED ON SYSTEM]* [PRECODED INTERVAL SIZE (2)] - [CALCULATED LOW VALUE])/NUMBER OF BUCKETS

Table 6-3. Default Values for Plots

ID #	Max Size of Plot	Lower Plot Limit	Upper Plot Limit
26	Unlimited	0.	456.
27	Unlimited	0.	456.
28	Unlimited	0.	114.
59	Unlimited	0.	*

* - The Upper Plot Limit for Plot ID #59 is determined dynamically by the program.

Limit = Memory configured on System * Histogram Interval Size for Histogram ID #18 that is originally coded into the program (set to a value of 2)

Table 6-4. Available Report Actions and Their (Default) Values (Part 1 of 2)

HISTG - Modify a histogram (see table 6-2)

PLOT - Modify a Plot (see table 6-3)

ON - Turn a specific report on - (all reports on except Memory Map and Out of Core Report)

OFF - Turn a specific report off - (all reports on except Memory Map and Out of Core Report)

TIME - Set a timespan(s) for reporting - (total time reported)

ALLOFF - Turn all reports off except those specified - (all reports on except Memory Map and Out of Core Report)

ALLON - Turn all reports on except those specified - (all reports on except Memory Map and Out of Core Report)

ERROR - Do not stop on an option request error - (stop on an input error)

DEBUG - Program debug requested - (no debug)

ALLOC - Stop program after a specified number of memory allocations have been requested - (entire tape processed)

NREC - Stop program after a specified number of tape records have been processed - (entire tape processed)

NOUSER - Do not print USERID on any report - (USERID printed on certain reports)

IDLE - Turn off all Idle Monitor reports - (all IDLE reports on)

WASTED,CORE,IO,CPU,RATIO,URG - Changes parameters used in the Excessive Resource Usage Report - (20K,50K,30MIN, 30MIN,5,40)

ABORT - SNUMBs not to report in the ABORT Report - (all SNUMBs that abort are reported)

PLTINT - Change Interval at which plots are printed - (10 MIN)

FSTSLV - Change the lowest allowable user program number - (14 decimal) must be expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If a user wants to request the time 4:07, he must input 0407. All times must include four characters.

If a start time, but no stop time, is desired, no characters should be entered after the minutes of the start time. If a stop time is requested, there must be a start time corresponding to it. If the user wants to start at the beginning of data collection and stop at some specified time, but is not sure of the start time, a start time of 0001 should be used. Figure 6-5 shows the format for this option.

6.1.11 Turn All Reports Off Except Those Specified (Action Code ALLOFF). All reports except those explicitly identified here are to be turned off. The inputs consist of

A B C . . . Y (max of 25)

where A through Y are the report ID numbers (table 6-1) to be turned on. The format is shown in figure 6-6. This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.12 Turn All Reports On Except Those Specified (Action Code ALLON). All reports except those explicitly identified here are to be turned on. The input consists of

A B C . . . Y (max of 25)

A through Y are the report ID numbers (table 6-1) to be turned off. The format is the same as action code ALLOFF (see figure 6-6). This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.13 Continue Data Reduction After an Input Option Error (Action Code ERROR). This code allows data reduction to continue when an error has been detected and reported in an input option request. The default value will abort data reduction and report the error. Only the Action Code card is required.

6.1.14 Debug For a Given Program Number (Action Code DEBUG). This is a debug option which supplies large amounts of output for a given program number. It should be used only in cases of data reduction problems. Card 1 contains the word DEBUG and card 2 contains a program number. A program number of -150 will provide detailed debug on system scheduler activities.

6.1.15 Stop After a Specified Number of Tape Records Processed (Action Code NREC). This option is useful when a tape problem occurs and the entire tape cannot be processed. When this occurs, the program will usually abort with an I/O error and some reports might be lost. If a tape error does occur during data reduction, the operator should type a "U" in

```
Card  1   A
      2   N M
      3   B C D E ...
```

where
A = The word TIME
N = Report NAME to be time spanned (table 6-1)
M = Number of different times appearing on Card 3. N and M must be separated by at least one blank.
B,C,D,E = Start and stop times used to define the time spans. Times must be separated by one or more blanks.

For each report that is to be timespanned, the entire sequence of three cards must be repeated.

Figure 6-5. TIME Action Code Format

```
Card  1 = A
      2 = N
      3 = B C D E ...
```

where
A = The word ALLOFF or ALLON
N = The number of report IDs appearing on card 3 cannot exceed 25
B,C,D,E = ID number of those reports not to be turned OFF/ON. All numbers must be separated by at least one blank.

Figure 6-6. ALLOFF/ALLON Action Code Format in order to stop data reduction processing prior to the tape error. The first card contains the word NREC and the second card contains the number of tape records to be processed.

6.1.16 Suppress USERID (Action Code NOUSER). This action code is used to suppress the printing of USERIDs on those reports where the USERID normally appears. Only the Action Card is required.

6.1.17 Turn Idle Reports Off (Action Code IDLE). This option will turn off histograms dealing with idle CPU information (i.e., report IDs 16, 19, 20, 21, 22, 31-34, 43-54, 55-58). The user should realize that these reports are useful in determining the I/O boundness of the system. However, on most systems, the idle trace is 70 percent of the entire tape, so that, by turning off this processing, processing time can be reduced by over 50 percent. Only the action card is required for this option.

6.1.18 Change Excessive Resource Limits Used in Excessive Resource Report (Action Code WASTED, CORE, IO, CPU, RATIO and URG). This report lists all jobs which are above a preset threshold for any of the following resources:

Wasted Memory
Excessive Memory
Excessive CPU time
Excessive I/O time
Excessive Ratio
Excessive Urgency

These limits are currently set to the values specified in table 6-4 and may be changed by using this option. The format for this option consists of Card 1 specifying the action code and Card 2 specifying the new threshold limit. This report is explained later in this chapter.

6.1.19 Eliminate SNUMBs From Abort Report (Action Code ABORT). This report lists all activities that fail to go to EOJ (i.e., Abort). The details of the report are given in section 6.3 of this chapter. At times, jobs are designed in such a way that they can be terminated only via a MME GEBORT or operator command. While these jobs do not go to EOJ, they have processed correctly and have not resulted in wasted computer resources. This option allows the user to request that these jobs not be included in the Abort Report. The first card contains the action code ABORT. The second card contains the number of jobs that will be deleted from the Abort Report. This number may not exceed 10. If more than 10 jobs are listed, only the first 10 will be deleted. The third card contains the SNUMB of each job to be deleted. Each SNUMB must be followed by at least one blank column.

6.1.20 Change the Plot Interval (Action Code PLTINT). Currently, all plots are outputted at 10-minute intervals. The plot interval controls the output of all plots; i.e., one plot cannot have a different time interval than another plot. The first card of this option contains the action code PLTINT. The second card contains the new plot interval inputted in minutes.

6.2 Processing

6.2.1 General. The reports of the MUDRP are intended to aid in the following:

- o System sizing - both memory sizing and processor utilization.
- o Job flow analysis - determining if and where a bottleneck exists and the user memory loading and the daily load distributions.
- o System perturbation measures - allows the user to evaluate how a new procedure or new load may alter the utilization of the system as well as determine the total utilization/capacity of the system.
- o Large user jobs - aid in identifying specific jobs which are misusing or "hogging" system resources.

Figure 6-7 illustrates how the monitor will pinpoint these various areas. For example, if the monitor indicates a large percentage of processor idleness with high memory demand and low memory availability, a dispatching or I/O bottleneck would be indicated. This would be caused by the I/O not completing its services in a sufficiently timely manner to allow full use of the processors. If processor use was very high and memory demand and availability were high, a memory allocation bottleneck or an overloaded processor would be indicated.

6.2.2 JCL. Figure 6-8 presents the JCL needed to run a total MUM reduction. The following points describe key features of the required JCL.

- o 74K required for memory
- o 15K sysout requirement would vary depending on amount of data collected. This figure would be significantly higher if the Memory Map or Out of Core Report were produced.
- o The DATA I* card is used to indicate the presence of data cards. All data cards must immediately follow this card. At least one data card must be present. That card will contain the word END and is used to signify the end of input data. The END card must be present even if no other data cards are desired. The data cards shown in the example are those recommended for most analyses.
- o An additional 8K will be required to load the MUM reduction program, but this 8K will be released immediately upon loading.
- o A JCL file is already established for the user under the file B29IDPXO/JCL/MUM.

```
$    IDENT         1820251/30/3044,C702
$    SELECT        B29IDPXO/OBJECT/MUM
$    TAPE          01,X1D,,18897
$    LIMITS        999,74K,-4K,15K
$    DATA          I*

(Following is a list of recommended data cards)

     ALLON
     10
     5 7 12 13 14 15 17 29 30 39 46 47
     SPECL
     3
     TS1 FTS TLNT
     END
```

Figure 6-8. JCL to RUN MUDRP

Table 6-5 shows all the MUDRP file codes and their corresponding reports.

6.3 Outputs

In this section, a simple explanation of how each report was derived from the data is given. Subsection 6.1 discussed how the ranges and other options of each report may be modified to fit an individual installation. While this section will provide some insight as to how an analyst should proceed to review all the reports produced by the MUDRP, section 14 provides a step-by-step approach as to how a memory analysis might be conducted.

Immediately prior to the output of the histograms, the user will find a printout containing processing information. Included in this information is the following:

- o Printout of all input options selected by user
- o Indication of multireel tapes that are being requested and have been mounted
- o Indication of the monitors that were active during data collection
- o Error messages - all error messages are either self-explanatory or else followed by the words "For Information Only." The latter messages are used by CCTC for future enhancements and as such can be ignored by the user.
- o If the time frame option was used, and indication of when the various time frames were reached.

6.3.1 MUM Title Page. The Memory Utilization Monitor (MUM) title page contains a summary of the systems configuration and activity over the measurement period (see figure 6-9). It displays the time the monitor was initiated and terminated, as well as identifying the system which was monitored and the tape number(s) containing the data. The configuration information is augmented by the amount of memory dedicated to the operating system itself, including that used by the memory allocation program. These figures will give the user a good idea of how much hard core space remains which could be used for SSA module hard core loading. If SSA cache is also configured, the amount of memory being used for this feature is also listed. The version number should be 09-83 CHG-7.

Immediately following is a summary of the work processed over the measurement period. The first set of lines provides information concerning the overhead generated by the actual data collection. The monitor name is given, its CPU time in seconds, an its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though

Table 6-5. File Code for MUM Reports (Part 1 of 2)

20	Activity Resource Report, Special Job Reports
21	IDENT Report
22	Special Job Report (temporary file)
23	Special Job Report (temporary file)
24	Urgency Over Time Report (temporary file)
26	Zero Urgency Job Report (temporary file)
27	Activity Abort Report
31	Plot 1 - (see table 6-1 for Plot Definition) (temporary file)
32	Plot 2 - (see table 6-1 for Plot Definition) (temporary file)
33	Plot 3 - (see table 6-1 for Plot Definition) (temporary file)
34	Excessive Resource Report
35	Plot 4 - (see table 6-1 for Plot Definition) (temporary file)
36	Used for outputting all plots
37	Used for outputting Out of Core Report, Memory Map, and Peripheral Allocator Report
42	Histograms, System Program Usage Report, Memory Statistics Report, Distribution of Urgency Over Time Report, Zero Urgency Job Report
45	Out of Core Report (temporary file)
51	Memory Map Report with one file required for each 128K Memory configured (temporary file)
52	Memory Map Report with one file required for each 128K Memory configured. (temporary file)
53	Memory Map Report with one file required for each 128K Memory configured (temporary file)
54	Memory Map Report with one file required for each 128K Memory configured (temporary file)

```
****** T H E  M E M O R Y  U T I L I Z A T I O N  M O N I T O R ******
                         VERSION 09-83  CHG-7

MONITORING ON 79-05-04 STARTED AT 11:39:50 AND COMPLETED AT 13:19:56 FOR A TOTAL TIME OF 1.67 HOURS
                      ON SYSTEM OSCC2  RUNNING 6.4.1D OF TAPE D0002

                      THE SYSTEMS CONFIGURATION CONSISTED OF:

                           2 - 6680 CENTRAL PROCESSORS
PROCESSOR OPTIONS - 2K CACHE, EIS OPTIONS INSTALLED, FAST MEMORY ACCESS, INSTRUCTION OVERLAP ON
                           2 - INPUT/OUTPUT MULTIPLEXORS
                               WITH 24 I/O CHANNELS
                         512 - 1024 WORD BLOCKS OF MEMORY
                          52 OF WHICH WERE USED BY THE HARD CORE SYSTEM ITSELF
                           3 OF WHICH WERE USED BY CALC
                          5K USED BY SSA CACHE (NOT HARD CORE)

            THE SYSTEM PROCESSED THE FOLLOWING OVER THE MEASUREMENT SESSION:

              MONITOR            TIME(SEC)            % OVERHEAD
                EXEC               100                   1.7
                MUM                 60                   1.0
                CPU                 40                    .7
                IDLE                40                    .7
                TOTAL                                    4.1

          175 ACTIVITIES WERE PROCESSED    AT A RATE/HOUR OF  104.89
              OF THESE 50 WERE SYSTEM SCHEDULAR ACTS
          125 ACTUAL ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF  74.85

          128 MOVES WERE PERFORMED         AT A RATE/HOUR OF   76.72
          266 SWAPS WERE PERFORMED         AT A RATE/HOUR OF  159.44
       201466 TIMES THE PROCESSORS WENT IDLE YIELDING A
            20 % IDLENESS OF THE PROCESSORS
         THE FIRST  PROCESSOR WENT IDLE 122479 TIMES, YIELDING 23 % IDLENESS
         THE SECOND PROCESSOR WENT IDLE  78987 TIMES, YIELDING 17 % IDLENESS

   THE MEMORY ALLOCATOR WAS CALLED    1907 TIMES -      0 OF WHICH RESULTED IN NO STATE CHANGE

THE TOTAL CPU TIME IN SECS WAS  6991  THE TOTAL IO TIME IN SECS WAS  15662  CPU/IO RATIO IS  0.446373
             WEIGHTED MEMORY SURPLUS IN K WORDS WAS    46
             WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 50  INCLUDES CALC AND PALC QUEUES
```

Figure 6-9. MUM Title Page Report - Idle Monitor Active listed separately it is also included under the monitor "MSM". The Monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM". These two special areas of code, within subroutine T7 (connect trace processing), are considered to be high usage areas and as such consume significant processing resources. In order to determine the true overhead of these areas, so that future code optimization can be considered, these areas are being reported separately.

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor, to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as

$$\frac{\text{(CPUTIME used by monitor)}}{\text{(Total Elapsed Time) x (number of Processors)}}$$

Following this are several lines describing the work performed during the monitoring session. These lines are self-explanatory.

If a termination record is not processed, either because the monitor aborted before a termination record could be written or else time frames were used, the lines describing GMC overhead will not be printed.

The number of times a processor went idle is derived from the idle processor traces captured by the IDLEM, with the percentage of processor idle also being gathered by the collection of idle state information. This is shown system-wide (i.e., for all the central processors and then individually for each processor). This information will not be present if the IDLEM was not active or if its output reports have been disabled by a data card option (see figure 6-10).

The number of memory allocator calls, as counted by the monitor, is shown. This much less than the number of calls to the multitude off SSA modules used by the Core Allocator and consists only of those that may have altered the memory state of the system. The second figure shows how many times a memory state change might have taken place and did not. This could be caused by no allocation being possible or by a call to the allocator pertaining to a matter other than allocation (i.e., a console message).

```
****** T H E  M E M O R Y  U T I L I Z A T I O N  M O N I T O R ******
                         VERSION 09-83  CHG-7

MONITORING ON 80-12-15 STARTED AT 12:39:46 AND COMPLETED AT 19:29:06 FOR A TOTAL TIME OF 6.82 HOURS
            ON SYSTEM NMCC2  RUNNING W64000 OF TAPE D0001

               THE SYSTEMS CONFIGURATION CONSISTED OF:

                  2 - 6680 CENTRAL PROCESSORS
PROCESSOR OPTIONS - 2K CACHE, EIS OPTIONS INSTALLED, FAST MEMORY ACCESS, INSTRUCTION OVERLAP ON
                  2 - INPUT/OUTPUT MULTIPLEXORS
                      WITH 16 I/O CHANNELS
                512 - 1024 WORD BLOCKS OF MEMORY
                 52 OF WHICH WERE USED BY THE HARD CORE SYSTEM ITSELF
                  3 OF WHICH WERE USED BY CALC
                 5K USED BY SSA CACHE (NOT HARD CORE)

        THE SYSTEM PROCESSED THE FOLLOWING OVER THE MEASUREMENT SESSION:

          MONITOR          TIME(SEC)          % OVERHEAD
            EXEC              300                1.2
            MUM               200                 .8
            TOTAL                                2.0

        1033 ACTIVITIES WERE PROCESSED     AT A RATE/HOUR OF  151.41
             OF THESE 33 WERE SYSTEM SCHEDULAR ACTS
        1000 ACTUAL ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 146.63

        1660 MOVES WERE PERFORMED          AT A RATE/HOUR OF  243.31
        6363 SWAPS WERE PERFORMED          AT A RATE/HOUR OF  932.66
                IDLE MONITOR WAS NOT ACTIVE OR WAS TURNED OFF

   THE MEMORY ALLOCATOR WAS CALLED   23409 TIMES -    14 OF WHICH RESULTED IN NO STATE CHANGE

THE TOTAL CPU TIME IN SECS WAS 31022  THE TOTAL IO TIME IN SECS WAS  61341  CPU/IO RATIO IS  0.505735
               WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 101 INCLUDES ONLY CALC QUEUE
               WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 150 INCLUDES CALC AND PALC QUEUES
```

Figure 6-10. MUM Title Page Report - Idle Monitor Off

The next line printed out is the Total CPU and I/O times in seconds and the ratio of CPU to I/O time. This figure gives the user an idea of whether the workload processed by the system is I/O or CPU dominant. It should be noted that these numbers are the amount of CPU and I/O time generated during the measurement period.

The next two lines give an indication of whether the system has a surplus or shortfall of memory. The weighted figure is calculated by using the following formula:

$$W = \frac{\sum_{i=1}^{N} \left(\begin{matrix}\text{memory} \\ \text{available}\end{matrix} - \begin{matrix}\text{demand for} \\ \text{memory}\end{matrix} \right) * \left(T_{i+1} - T_i \right)}{\text{TOTAL TIME}}$$

Where i = calls to the core allocator

T   T - length of time over which memory availability was in this state.

If W comes out positive, there is a core surplus and if W comes out negative, there is a core shortfall. In the first line, the demand for memory is taken only from the Core Allocator's queue. In the second line the demand for memory is taken from the demand in both the Core Allocator and Pripheral Allocator queues. The Peripheral Allocator's queue consists of the memory demand that is currently being processed by the Peripheral Allocator and has not yet reached the Core Allocator. The Peripheral Alloctor will stop transferring jobs to the Core Allocator when the Core Allocator's queue reaches a predefined length. This second figure presents a truer picture of memory availability. Jobs from the Peripheral Allocator are only included if they have been completely processed by the Peripheral Allocator. These figures present a good first indication of whether or not availability of memory is a system constraint. In calculating demand, a job is only included if it does not have a zero urgency. Any activity with a zero urgency will not be considered to have a core demand unless the activity is in a loading (activity 0) or terminating status.

6.3.2 System Program Usage. The report immediately following the title page provides an overview of the system program load on the memory subsystem. The data presented consists of the following:

- $\frac{\text{Total Memory Time for This System Program * 100}}{\text{Memory Time for all Programs}}$

This figure would indicate what percentage of the total memory time was used by this program.

- Percentage of the Elapsed Time in Memory

- $\frac{\text{Total Size Time Product for This System Program * 100}}{\text{Total Size Time Product for all Programs}}$

This figure would indicate what percentage of the total size time product was used by this program. The size-time product of a job is an attempt to determine the memory effect of a job based not only on its size, but on the length of time that it runs. A 20K job that runs for three hours might be more detrimental to a system than a 60K job that runs ten minutes.

- $$\frac{\text{Total Size Time Product for This System Program} * 100}{\text{Total Size Time Produce Available to System}}$$

Where Total Size Time Produce Available to System = (The Elapsed Run Time) * (Total Allocatable Memory)

The next two figures are weighted memory sizes for this program. The first figure is the weighted memory size of this program while it is in memory. Therefore, if TSS was in memory during three different time periods for 1/2 hour, 3/4 hour, and 1 hour, and during these periods its memory size was 40K, 100K, 180K respectively, its weighted in memory size would be calculated as follows:

$$\text{Weighted (IN)} = \frac{(40)*(.5)+(100)*(.75)+(180)*(1)}{2.25}$$

$$= \frac{275}{2.25} = 122\text{K}$$

Had the calculation not been weighted by time, the average size of TSS would have been:

$$\frac{(40)+(100)+(180)}{3} = 73\text{K}$$

In the above calculation, the report would be stating that the amount of memory being taken away from the system, by TSS, was 122K. However, what if TSS was swapped for 50 percent of the total elapsed time? Then TSS really did not take 122K from the system, but rather only 61K. The second weighted figure takes into account the total time the program was actually in memory.

The final figure is the number of times this program was swapped.

In addition to the standard system programs, any jobs requested by the user, to be considered as system jobs, also appear in this report. In figure 6-11 we see 6 user-requested jobs appearing on the report. The user had actually requested nine jobs to be considered as system jobs, but three of those jobs never appeared. In a system using multicopies of TSS only TS1 (prog #5) will appear in this report. Other copies of TSS must be requested by user input option "MASTER". In a WWMCCS system, a program is considered to be a system program if it has a program number less than decimal 14. Commercial users should use the FSTSLV option to change the

THE SYSTEM PROGRAM USAGE OF MEMORY WAS:

PROGRAM	% USED OF MEMORY USE TIME	% OF TIME IN MEMORY	% USED OF SIZE-TIME USED OF MEMORY	% USED OF SIZE-TIME AVAILABLE	WEIGHTED IN/TTM	# SWAPS
(001) GMF MONITOR	7%	100%	5%	5%	23/ 24	0
(002) PERIPHERAL ALLOCATOR	5%	71%	2%	2%	17/ 12	988
(003) GENERAL SYSTEM OUTPUT	6%	93%	6%	5%	26/ 25	58
(004) REMOTE INPUT	2%	35%	1%	1%	16/ 5	415
(005) TIME-SHARING EXECUTIVE	7%	100%	31%	28%	129/129	0
(006) TEST AND DIAGNOSTICS	0%	1%	0%	0%	6/ 0	0
(010) LOG ON - SECURITY	2%	38%	1%	0%	11/ 4	292
(011) FILE MANAGEMENT SYSTEM	0%	12%	0%	0%	9/ 1	0
(012) WW NETWORK CONTROL (NCP)	6%	88%	5%	5%	27/ 24	79
(013) WW TELNET	4%	65%	2%	2%	16/ 10	150
(014) WW T2EXEC	3%	47%	3%	3%	28/ 13	233
(015) WW TLCF	2%	38%	1%	1%	16/ 6	54
(016) DMS PROCEDURES (DMTEX)	0%	0%	0%	0%	4/ 0	1
SPEC 1 HEALS	1%	21%	0%	0%	7/ 1	150
SPEC 4 NACE	3%	50%	2%	2%	18/ 9	1652
SPEC 6 VIDEO	7%	99%	1%	1%	6/ 5	2
SPEC 7 LGEN	7%	99%	4%	3%	17/ 16	4
SPEC 8 NEWRM	0%	8%	0%	0%	3/ 0	0
SPEC 9 DMSTA	5%	79%	1%	1%	8/ 6	717
				TOTAL	387/290	

Figure 6-11. System Program Load

value of 14 to an 8. In addition, the GMC Monitor, $HEALS, VIDEO and $TRAX are all considered to be system programs.

6.3.3 MUM Reports. The following paragraphs describe the reports output by MUM.

Report numbers 1-50 are all presented in histogram format (see figure 6-12). At the top of the report, the system name, as well as the time and date of data collection, are given. This is followed by the title line of the histogram. Column number 1 indicates the number of occurrences of a given event, with column number 5 describing the event. In figure 6-12, we find that 229 times there were 0 user activities in memory, while 2899 times there were 5 activities in memory. Column number 2 is simply a running total of column number 1. Therefore, the second line in column number 2 (2008) is merely a running total of the first two lines of column number 1 (229 + 1779). The fourth column is the percentage of all activities which will fall into that line of the report. For example, 4063 entries out of a total 23410 entries indicate 2 activities in memory. This results in a 17.35 percentage figure. This means that 17.35% of all measurements (4063/23410) showed 2 activities being in memory. Column number 3 is simply a running total of column number 4. It presents the percentage of measurements which will fall into a given line, or earlier line. For example, 25.93% of all measurement showed the number of activities in memory to be 2 or less. There is a graphic display of these measurements presented to the right of the fifth column. At the bottom of the report, summary information is provided and is calculated in the standard statistical manner.

In figure 6-13, we see a similar histogram report. As displayed by column 5, we find that each line of the histogram represents a range of values, with an interval size of 200. This interval size can be modified by the user. The lowest value in this histogram is 0 (modifiable by the user) and the size of the histogram defaults to 45 lines (also modifiable by the user). Actually, for this run, the lowest value recorded was 42. Since we can output only 45 lines and each line represents a range of values of 200, the largest value that could be reported would be 9,000 (200 x 45). If a measurement falls outside this maximum value, it is reported as an out-of-range value. In figure 6-13, we find the 21 measurements exceeded 9,000. The average of these 21 measurements was 20188.48. The first line of the summary includes all measurements that were taken. Therefore, 21 out of 79 entries (26 percent) of all measurements were out of range. The average of all measurements taken was 5953.62, while the average of the in-range masurements, (all out of range values are eliminated) was 799.62.

6.3.3.1 Report 1 - Memory Demand Sizes of New Activities in 1K Word Blocks. This report shows the demand size, in 1K-word blocks, of each individual user activity as it was first seen by the memory allocator. The demand sizes are presented to the allocator by the Peripheral Allocator. This is a good measure of the memory demand load of a systems

```
                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2  AT 12:39:46 ON 80-12-15

                              THE # OF USER ACTIVITIES IN MEMORY

INDIV.  CUMUL.  CUMUL.   INDIV.  NUMBER             PERCENT OF OCCURRENCE                    REPORT
NUMBER  NUMBER   PRC      PRC     WAIT    00  05  10  15  20  25  30  35  40  45  50           15
                                          I....I....I....I....I....I....I....I....I....I....I
  229     229    0.978    0.978   0-   0  IX
 1779    2008    8.578    7.599   1-   1  IXXXXXXXX
 4063    6071   25.933   17.356   2-   2  IXXXXXXXXXXXXXXXXX
 6597   12668   54.114   28.180   3-   3  IXXXXXXXXXXXXXXXXXXXXXXXXXXXX
 7062   19730   84.280   30.167   4-   4  IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
 2899   22629   96.664   12.384   5-   5  IXXXXXXXXXXXX
  633   23262   99.368    2.704   6-   6  IXXX
  134   23396   99.940    0.572   7-   7  IX
   14   23410  100.000    0.060   8-   8  I

23410 ENTRIES TOTAL    AVERAGE =  3.30145    VARIANCE =  1.622    STANDARD DEVIATION =  1.273
```

Figure 6-12. Standard Histogram Report

```
                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:50:03 ON 81-03-13

                       THE ELAPSED DURATION OF USER ACT IN TENTHS OF A SECOND

INDIV.  CUMUL.  CUMUL.  INDIV.   TENTHS              PERCENT OF OCCURRENCE                     REPORT
NUMBER  NUMBER   PRC     PRC     SECOND   00   05   10   15   20   25   30   35   40   45   50      20
                                          I....I....I....I....I....I....I....I....I....I....I
  22      22    27.848  27.848    42- 199 IXXXXXXXXXXXXXXXXXXXXXXXXXXXX
  14      36    45.570  17.722   200- 399 IXXXXXXXXXXXXXXXXXX
   5      41    51.899   6.329   400- 599 IXXXXXX
   2      43    54.430  .2.532   600- 799 IXXX
   1      44    55.696   1.266   800- 999 IX
   0      44    55.696   0.     1000-1199 I
   1      45    56.962   1.266  1200-1399 IX
   2      47    59.494   2.532  1400-1599 IXXX
   1      48    60.759   1.266  1600-1799 IX
   0      48    60.759   0.     1800-1999 I
   1      49    62.025   1.266  2000-2199 IX
   3      52    65.823   3.797  2200-2399 IXXXX
   0      52    65.823   0.     2400-2599 I
   0      52    65.823   0.     2600-2799 I
   1      53    67.089   1.266  2800-2999 IX
   1      54    68.354   1.266  3000-3199 IX
   1      55    69.620   1.266  3200-3399 IX
   0      55    69.620   0.     3400-3599 I
   2      57    72.152   2.532  3600-3799 IXXX
   0      57    72.152   0.     3800-3999 I
   0      57    72.152   0.     4000-4199 I
   1      58    73.418   1.266  4200-4399 IX

 79 ENTRIES TOTAL    AVERAGE = 5953.62024    VARIANCE = 04446990.000    STANDARD DEVIATION = 10219.931
          21(26%) OUT OF RANGE    AVERAGE FOR THESE =     20188.48    IN RANGE AVERAGE =     799.62
```

Figure 6-13. *Out-of-Range Histogram*

operation and can be used to set System Schedular classes to correctly balance the load cross varying memory size jobs. This report shows the percentage of activites which had a particular memory size. For this report, an entry is made for each new user activity demand at each allocator call. See Report 10 for an explanation of user vs. system activity.

6.3.3.2 Report 2 - The Memory Demand Size of All Demand Types. This report contains the information in Report 1, with the addition of all other individual demand types. These include activities that are swapped or involved in a memory compaction procedure. This report should be similar to Report 1, unless a great amount of GEMORE, GERLEC, or swap operations are performed by the users load. This would alter the memory size demands from that seen by the allocator at the initial request. For this report, an entry is made for each activity with an outstanding demand for each allocator call. Activities with an urgency of 0 are not counted.

6.3.3.3 Report 3 - The Potential Time-Weighted Memory Demand. This report shows the total memory required if all jobs currently in the system were to have their memory demand satisfied at one time. The data shown in this report is the sum of all jobs currently in memory, plus all jobs currently swapped, plus all jobs waiting for their original allocation of memory, plus all jobs currently in the system with a zero urgency. This report is an indication of the worst case demand that may be potentially placed upon the system. An entry is made to this report at each allocator call. See report 16 for an explanation of time weighting.

6.3.3.4 Report 4 - The Demand That Was Outstanding When a Processor Went Idle. This report shows the sum of demand for all activities in the system including outstanding GEMOREs. It is a distribution of memory demand that is not satisfied, across the measurement session. It should be remembered that all data is collected at the Core Allocator and does not represent the full system load. Portions of the load may be held in the System Schedular and the Peripheral Allocator. Activities with an urgency of 0 are not counted. An entry is made only if a processor has gone idle since the last allocator call. If a large demand should be outstanding during processor idleness, a system bottleneck may be present. In this case, memory is probably fully utilized (i.e., demand cannot be satisfied), but the activities that are occupying memory are not using the processor, (i.e., a processor has gone idle). This is a good sign of an I/O backlog. IDLEM data is used to produce this report. If the Idle Monitor was not active, this report will not be produced.

6.3.3.5 Report 5 - The Total Amount of Available Memory. The total amount of available memory is a key indicator of the system memory utilization. If this amount is continually low, the memory is being fully utilized and possibly in need of expansion. A continually high amount may indicate another system bottleneck or an excess of memory. This report, when used in conjunction with Reports 3, 4, and 6 should give a good first-level indication of system memory utilization. It should be noted that the availability shown here exists in all quadrants. The availability is the sum of any and all "holes" in the system and does not mean that this memory is contiguously available.

Any activity with an urgency of 0 that is currently in memory will have its memory size included in this availability figure. The reason for this is that if memory becomes a constraint, these activities can be swapped and their memory will become available for use.

For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 8 provides a more statistically accurate representation of this data.

6.3.3.6 Report 6 - The Memory Available When a Processor Went Idle. The previous report is repeated with the additional restraint that a processor has gone idle since the last allocator call. This aids in identifying either a bottleneck or a lightly loaded system.

For this report, an entry is made at each allocator call that had a processor go idle since the last allocator call. IDLEM data is used to produce this report. This report will not be produced if IDLEM was not active or the IDLEM Reports have been disabled via user input command.

6.3.3.7 Report 7 - The Time-Corrected Total Demand Outstanding. See report 16 for an explanation of time correction. The time-corrected total demand is the sum of all requests for memory known to the allocator as indicated in report 4. Activities with urgency 0 are not counted.

6.3.3.8 Report 8 - The Time-Corrected Memory Available. See report 16 for an explanation of time correction. This report reflects the time-corrected amount of total memory available as indicated in report 5.

6.3.3.9 Report 9 - The Number of Activities Waiting for Memory in Allocator Queue. This report identifies the depth of the allocator demand queue and includes all activites that are waiting for memory allocation. Activities with a 0 urgency are not considered as waiting for memory. This report aids in determining if too many or too few activities are getting to the Core Allocator from the Peripheral Allocator. For this report, an entry is made at each allocator call. For most analyses, this report will not be used since report 11 provides a more statistically accurate representation of this data.

6.3.3.10 Report 10 - The Number of User Activities Waiting Memory in Allocator Queue. This report is the same as report 9 except that it only counts those activites of a slave job as identified by their program number (program number 14 or greater). In order to change this program number test, the user should see Input Action FSTSLV. In addition, the user may specify up to ten additional programs that he wants considered as system programs, even though their program number exceeds 14. The user should see Input Action MASTER in order to select this option. This report indicates the "user" work waiting allocation. For this report, an entry is made on each allocator call. For most analyses, this report will not be used since report 12 provides a more statistically accurate representation of this data.

6.3.3.11 Report 11 - The Time-Corrected Number of Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of activites waiting memory as in report 9.

6.3.3.12 Report 12 - The Time-Corrected Number of User Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of user jobs waiting memory in the allocators queue as in report 10. See report 10 for additional user options.

6.3.3.13 Report 13 - The Number of Activities Waiting Memory When a Processor Went Idle. Report 9 is the basis for this report, with the additional criteria that a processor must have gone idle since the last allocator call. An entry is made for each allocation where a processor has gone idle since the last call. IDLEM data is used to produce this report. This report will not be produced if IDLEM is not active or the [illegible] reports were disabled via user input commands.

6.3.3.14 Report 14 - The Number of Activites Residing in Memory. This report represents the number of activities allocated memory. It indicates the multiprogramming depth the system is obtaining. It is probably an upper level since an activity is allocated memory prior to and past actual usage. Any activity in memory, with a 0 urgency, is not considered as residing in memory. For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 16 provides a more statistically accurate representation of this data.

6.3.3.15 Report 15 - The Number of User Activities in Memory. The activites shown in this report are those that are in memory and have a program number greater than or equal to 14. These are user programs. For this report, an entry is made at each allocator call. As explained in report 14, any activity which has an urgency of zero will not be counted as being in memory. See report 10 for additional user options in defining system jobs and user jobs. For most analyses, this report will not be used since report 17 provides a more statistically accurate representation of this data.

6.3.3.16 Report 16 - The Time-Corrected Number of Activities in Memory. This report presents the same information as in report 14. The number of entries at each allocator call is determined by the time since the last allocator call. The result is a simulation of a uniform sample rate of allocator calls. Therefore, the noncorrected reports display the distributions as seen by the allocator itself. The time-corrected reports present the time weighted distributions. As an example assume that three

IDLEM data is used to produce these reports. These reports will not be produced if IDLEM was not active or if the IDLEM reports have been disabled by user input command.

6.3.3.39 Report 50 - Original Allocation Time for User Memory in I/O Second. This report gives the time each user activty waited for its original allocation of memory. See report 10 for an explanation of user and system activities.

6.3.3.40 Report 51 - The Time-Corrected Percent of Assigned Memory Used. This report gives the time-corrected percentage of slave memory used over the monitoring period. Any memory being utilized by jobs with zero urgency will not be included in the memory-used figure for this report. See report 16 for a definition of Time Correction.

6.3.4 Activity Resource Usage Report. For each activity known to the monitor, a detailed Resource Usage Report is made upon termination of the activity. The report is ordered by termination time sequence, and the resource usage is that known to the system at the last allocator call (refer to figure 6-15).

Each activity is displayed via the SNUMB and activity number followed by the CP an I/O charge times expressed in milliseconds. This is the CP and I/O times generated during the monitoring session. The size-time product is the total K words times the microseconds of allocation time, which gives a better expression for the memory used by the job than the size of the job. The minimum and maximum core requirements of the job are then shown, including the activity Slave Service Areas (SSAs) as well as slave size.

The elapsed time, in hours, an activity was known to the allocator is followed by the number of times the job size changed for any reason. The wasted core column is calculated from the job Slave Prefix Area (SPA) word 37 octal. This is filled by the System Loader and may not be valid for all job types (i.e., an H* file is not loaded in the normal system load manner). This column is shown in order to help locate users that do not have the $LIMITS card set correctly for the memory being used. If the user appears to be requesting excessive core on his $LIMITS card, he may be using this extra space as a spare buffer area. If this figure shows an excessive misuse of the $LIMITS card, the user should be contacted and questioned.

The next two columns provide a count of the total number of swaps and moves incurred by the activity. The final columns of the entry gives memory allocation time, wait time, swap time, memory time, and GEWAKE time, all in tenths of a second for each activity. An entry will be made in this report for every activity of a job, when the activity completes. Upon termination of the monitor, the resource usage of all actiites known to the allocator will be reported, including system jobs. This output follows a full line of asterisks to denote that no termination records were found for these activities.

COLLECTED ON SYSTEM NMCC2 ON 80-12-15 AT TIME 12:39

ACTIVITY RESOURCE USAGE REPORT - REPORTED PER ACT SIZE

SNUMB-ACT		CPU & IO TIME	(MS)	SIZE-TIME PROD	MIN	MAX	ELAPSED TIME	SIZE CHANGE	WASTED CORE	SWAPS	MOVES	TIME(.1 SEC) SPENT IN ALLOC	SWAP	MEMORY	GEWAKE
7332T-	0	44	108	1.5829E 08	31	31	0.003	0	0	0	0	51	0	55	0
7329T-	0	12077	29043	2.6766E 09	32	32	0.023	0	0	0	0	0	0	837	0
$GENB-	0	275	1003	1.7747E 07	11	11	0.001	0	0	0	0	2	0	16	0
7339T-	0	1300	2483	1.5904E 08	40	40	0.008	0	0	0	0	234	0	41	0
7348T-	0	1421	2646	1.3063E 08	31	31	0.042	0	0	0	0	1463	0	43	0
7338T-	1	991	2761	8.1595E 07	12	12	0.016	0	2	0	0	512	0	67	0
7338T-	2	2267	4284	2.6228E 08	25	36	0.003	2	0	0	2	15	0	83	0
7338T-	3	2761	4297	3.7449E 08	45	56	0.006	-2	0	0	2	131	0	83	0
7354T-	0	1360	2909	2.1781E 08	40	40	0.002	0	0	0	0	5	0	55	0
7338T-	4	1766	3848	3.5940E 08	42	53	0.005	1	6	0	1	110	0	76	0
XXXXX-	0	2388	13991	4.8413E 08	11	11	0.016	0	0	0	0	146	0	440	0

FOLLOWING INFO MAY BE INCOMPLETE DUE TO LOST, DATA, NO EOF, OR ACTIVITY WAS ACTIVE WHEN MONITOR ENDED.

SYSTEM SCHEDULER CPU TIME, IO TIME AND SIZE-TIME PRODUCT ARE 43122911 129261145 0.19D 12

$PALC-	1	615091	2394913	2.9917E 11	17	17	6.822	0	0	988	108	0	17546	175982	52077
SYSOT-	0	626346	21551415	6.2101E 11	27	27	6.822	0	0	58	119	0	1877	230004	13728
$RGIN-	0	52768	167826	1.3920E 11	16	16	6.822	0	0	415	28	0	34039	87000	124567
-TSS—	1	12163737	13709653	3.1906E 12	77	187	6.822	58	0	0	0	0	0	245610	0
LOGON-	0	91433	59898	1.0459E 11	11	11	6.822	0	0	292	23	0	10302	95078	140228
FSYS -	0	0	90	2.7697E 10	9	9	6.822	0	0	0	0	0	0	30774	214835
NCP -	1	387705	72254	5.9289E 11	26	28	6.822	2	0	79	43	0	3452	216492	25664
TELNE-	1	231955	27456	2.5972E 11	16	17	6.822	2	0	150	22	0	5644	161220	78745
FTS -	1	216277	488796	1.7443E 11	24	51	4.732	176	0	191	111	9	3053	55649	111637
TLCF -	1	35090	32496	1.5089E 11	16	16	6.822	0	0	54	44	0	1839	94307	149462
DMTEX-	1	0	60	7.6714E 06	4	4	6.822	0	0	1	0	0	13	19	245577
DMSTA-	1	450328	1113536	1.5680E 11	8	8	6.822	0	0	717	76	0	13430	195994	36182
VIDEO-	1	116630	496	1.4735E 11	6	6	6.822	0	0	2	0	0	21	245588	0
7871T-	2	505393	384600	6.4366E 10	38	38	0.645	0	0	45	9	26	6263	16938	0
WAITL-	999	0	0	0.	58	58	0.175	0	0	0	0	6291	0	0	0
$GENB-	0	673795	2019705	1.9096E 11	** MEMORY USE TIME * 8.6242E 09										

Figure 6-15. Activity Resource Usage Report

```
                              MEMORY MAP OF    THE SECOND HALF OF THE SECOND QUADRANT                               PAGE    1

                400       410       420       430       440       450       460       470       480       490       500       510
LINE 8901234567890123456789012345678901234567890123456789012345678901234567890123456789012345678901234567890123456789012
  1- -----------2930T-003U32-------------------------*    *--NACE  -001U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  2- -----------2930T-003U32-------------------------*    *--NACE  -001U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  3- -----------2930T-003U32-------------------------*    *--NACE  -001U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  4- -----------2930T-003U32-------------------------*                      *---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  5- -----------2930T-003U32-------------------------**----SPALC-001U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  6- -----------2930T-003U32-------------------------**----SPALC-001U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  7- -----------2930T-003U32-------------------------**----SPALC-001U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  8- -----------2930T-003U32-------------------------**----SPALC-301U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
  9- -----------2930T-003U32-------------------------**----SPALC-301U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 10- -----------2930T-003U32-------------------------**----SPALC-301U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 11- -----------2930T-003U32-------------------------**----SPALC-301U00----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 12- -----------2930T-003U32-------------------------*                      *---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 13- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 14- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 15- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 16- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 17- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 18- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 19- -----------2930T-003U32-------------------------*    *--SHEAL-000U00--**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 20- -----------2930T-003U32-------------------------*                      *---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 21- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 22- -----------2930T-003U32-------------------------**----SPALC-801U62----**---------$SYOT-000U62----------*VIDEO*SSA CACHE CO
 23- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 24- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 25- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 26- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-000U62----------*VIDEO*SSA CACHE CO
 27- -----------2930T-003U32-------------------------**----SPALC-801U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 28- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 29- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 30- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 31- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 32- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 33- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-000U62----------*VIDEO*SSA CACHE CO
 34- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-000U62----------*VIDEO*SSA CACHE CO
 35- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 36- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-000U62----------*VIDEO*SSA CACHE CO
 37- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 38- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 39- -----------2930T-003U32-------------------------**----SPALC-001U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 40- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 41- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 42- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 43- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 44- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 45- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CC
 46- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 47- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 48- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 49- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
 50- -----------2930T-003U32-------------------------**----SPALC-301U62----**---------$SYOT-U00U62----------*VIDEO*SSA CACHE CO
```

Figure 6-17. (Part 3 of 3)

-JOB01-XXXU05-	(16K) - including SSAs
*JOB02-XXX--	(12K) - urgency discarded
JOB03	(7K) - activity discarded
JOB04	(5K) - asterisks discarded
****	(4K) - no identification

As can be seen in figure 6-17, Part 1, every line of the figure has a line number ranging from 1-50. In addition, there is a page number in the upper right corner. When the user wants to match this picture of the first half of the first quadrant with its corresponding half of some other quadrant, the following steps should be followed:

1- Match page numbers (see figure 6-17, Part 2, Part 3)
2- Match the line numbers from identical page numbers.

Two special names can appear in the memory map. If SSA cache memory is configured the following letters will be found in the map, depending on the size of the SSA cache memory:

*SSA CACHE**	12K	
*SSA CACHE	10K	
*SSA CAC	8K	
SSA C	5K	(see figure 6-17, Part 3)

If memory has been released from the system, then the letters *-RELEASED* will appear in the map. This will be repeated depending upon how much memory has been released.

6.3.7 Demand List Report. The Demand List Report shows the memory demand outstanding for each memory state displayed on the memory map. The correlation is made using the same line numbers as the half quadrants of the maps themselves (refer to figure 6-18).

The Demand List Report shows the total memory available, the number of jobs waiting memory, the demand request sizes for each job waiting memory. The memory available is the sum of all holes in memory.

6.3.8 Activity Abort Request. This report is directly related to the Activity Resource Usage Report. This report is produced whenever the Activity Report is produced. For every activity that aborts during the monitoring session, an entry is made to this report. The entry gives the SNUMB, Activity Number, Abort Code, CPU Time, Run Hours, USERID, and IDENT for the activity.

The Abort Code is either an octal number or an alphanumeric value. The meaning of these codes can be found an Appendix A of Honeywell Manual DD19 (GCOS).

statement that the $LIMITS card appears to be requesting more memory than is actually required by this job. The user should be questioned in order to determine if this is in fact true. In the Honeywell System, a user will receive whatever amount of memory requested on the $LIMITS card, whether or not the amount of memory is actually needed. The Ratio column shows the ratio of the total elapsed time for an activity divided by the total memory time for the activity. This value gives an indication of the activity lengthening factor; i.e., how run time is affected by resource contention. For those activities using excessive memory, the report also indicates, under the MEM MIN column, the amount of time the activity was in memory. The value being used for the urgency check is the average urgency recorded for the activity and not the maximum urgency of the activity. The default values for an entry being made to this report are listed in table 6-4. These values can be changed via a previously described input option. This report will be produced whenever the Activity Resource Report is produced and will be turned off whenever the Activity Resource Report is off (see figure 6-21).

6.3.11 Allocation Status Report. This report will track an activity as it proceeds through different phases of Allocation. The report lists the SNUMB-Activity #, amount of memory the activity required, its current status, the time it entered that phase of allocation, the time it completed that phase of allocation, the total time spent in a given phase of allocation, the device type, if any, that it was waiting for, and the number of devices, if any, the activity was waiting for. Due to the manner in which data is collected for this report, it is possible that certain phases of allocation will be missed, especially if that phase of allocation occurs within a short time span. This report will give a good indication of how long it is taking activities to pass through the various allocation phases prior to core allocation. Following is a list of the more common phases of allocation and their meanings:

New Act - Activity has just entered the Peripheral Allocator
Wait Media - Activity is waiting for a device
Wait Mnt - Activity is waiting for a disk pack or tape to be mounted
Core Queue Full - Activity has been completely processed and is waiting for the Peripheral Allocator to send the job to the core allocator
Alloc Done - Activity has been sent to core allocator. For this case the stop time and total time columns have no real meaning. These columns simple are reporting the amount of time it took the monitor to realize that the activity had reached the core allocator
LIMBO - Activity is in Limbo and has not even been granted permission to run
HOLD - Activity is in Hold and has not even been given permission to run
SCHED - Activity was in one of the System Scheduler queues.

Only activities found to be in a state for more than 600 seconds will be reported. This limit can be changed by using the PALC input option. See Figure 6-22 for a sample of this report.

[illegible] 12 Plot Reports. Four different plot reports are produced by the data reduction program. All plots are produced under 10-minute intervals, where the interval can be modified by the user. At every allocator call, the various parameters to the plots are accumulated and every 10 minutes, the accumulated parameters are averaged and an average value is output

7.2 Data Collection Methodology

The MSM in the General Monitor Collector processes GCOS trace types 7 and 15 and collects information to monitor the usage of the entire disk subsystem. The information collected on the occurrence of the above traces enables the MSMDRP to identify the activity issuing the I/O request, the file being accessed, the disk pack upon which the file is located, arm movement required in order to accomplish the requested file accessing, and the type of accessing being requested; e.g., read, write, write verify, etc.

If the system being monitored by the MSM is configured with SSA Cache Core, the MSM will create two direct transfer traces (types 73 and 76) in order to collect data to analyze the effectiveness of SSA Cache Core. The method for generating these new direct transfer traces is described in subsection 5.2.2, and the formats for the MSM generated records used by the MSMDRP are described in subsection 5.4.3.

Finally, if the system being monitored by the MSM is configured with FMS Catalog Cache or is utilizing disk in core space tables, a data record is generated so that the MSMDRP can report on the effectiveness of FMS Catalog Cache or in core space table buffering.

7.3 Analytical Methodology

An evaluation of the Mass Storage Subsystem reports produced by the MSMDRP requires concurrent use of the reports produced by the Channel Monitor Data Reduction Program (CMDRP). Chapter 14 provides a detailed description of the procedure to be followed in such an evaluation. Subsection 8.3 provides a detailed description of the entire I/O process, and the traces generated during the processing of an I/O request. In general, the CMDRP is used to identify channels and/or devices which are acting as bottlenecks to the efficient operation of the system, while the MSMDRP reports are used to determine the exact activities, files, and file codes that are causing the contention uncovered by the CMDRP reports. The MSMDRP reports will also identify those devices experiencing seek elongation problems and the files upon these devices which are responsible for the seek elongation. Finally, the MSMDRP reports will identify those files that are candidates for device relocation or placement into Hard Core or SSA Cache Buffer space.

Before a user conducts a Mass Storage Subsystem Evaluation, it is important to have an understanding of the entire I/O process. Subsection 8.3 provides a detailed description of the entire I/O process and all traces generated during the processing of an I/O request. In this subsection, a description of only the connect (trace type 7) event will be presented.

Each time a system program or application program issues an I/O request (read disk/tape, write disk/tape, seek, etc. . .) the GCOS system will generate a trace type 7 (connect event). Upon the occurrence of this event, several internal tables are updated and it is these tables that the

MSM references in order to generate its data record. A program's SSA area contains tables for the Peripheral Assignment Table (PAT) and the PAT Pointer. These are used to describe the device and space allocation for a particular file and the file code to correlate a user file code to the PAT and the device on which that file is allocated. The .CRIO and .CRCT tables contain descriptive information concerning device and channel configuration. Finally, the program's SSA area also contains an area which is used for I/O entries. These entries are each 11 words long and contain detailed information concerning the I/O just requested. They are referred to as the 11 word I/O queue entry.

I/O requests can be of two types (single or multicommand). Multicommands are of the type seek-read, seek-write, or seek-write verify. Single commands can be status requests of certain types, or reads/writes, where seeks are not required. These different types of I/O commands are processed and reported in different fashions by the various MSMDRP reports (see individual output reports). Finally, whenever the system generates a multi I/O command, it is necessary for the system to record the actual seek address being requested. Normally, this seek address is stored in I/O queue word number 4. Whenever the MSM processes a multicommand, it expects to find a valid seek address at this location. However, there are certain occurrences when a multicommand is issued and I/O queue word 4 does not contain a valid seek address. In these cases, Bit 32 of I/O queue word 2 is set to a 0.

7.4 Data Reduction Methodology

The MSMDRP currently uses random I/O (File 58) to process histogram data for the Device Space Utilization and Device Seek Movement reports. This feature allows the MSMDRP to process an unlimited number of devices with a minor increase in memory requirements. As delivered, the MSMDRP will process data describing 75 mass storage devices and 40 mass storage channels. It will produce 130 unique histograms with no random I/O. If the number of channels or devices is insufficient, the user will need to edit file B29IDPXO/SOURCE/MSM. The user should enter the edit subsystem and process the following command:

```
B RS:/NRDEVXX=XX75/;*:/NRDEVXX=XX new number of devices/

B RS:/NRCHANXX=XX40/;*:/NRCHANXX=XX number of new channels/
```

For each additional device, the size of the program will increase by 10 words and for each additional channel, the program will increase by 45 words. For the above edit, the character "X" signifies a space.

The next variable that will need to be changed is RPTCNT. This number represents the total number of histograms and reports that will be processed with no random I/O. To calculate the value required, the following formula should be used.

(number of devices actually configured)*2+8

If this value is less than 138 (65 disk devices), no change is required. If the value required is greater than 138, the user may alter this value. This will help to limit the amount of random I/O being performed but will increase storage by 80 words for each increment above 138. This trade-off between CPU/IO time and memory must be made at the discretion of the user. In order to change this value, the following edit function should be performed:

```
B RS:/RPTCNTX=X138/;*:/RPTCNTX=Xnew value/
```

As in the earlier edit example, the character "X" should not be typed, but is being used to represent a blank column.

After performing the above edits, the user should recompile the source program by entering the card subsystem and issuing a run command.

7.5 MSMDRP Output

The MSMDRP produces a series of 24 reports listed in table 7-1 over which the user has limited control. Those eight reports with a NAME=codename designation offer greater parameter control to the user. This parameter control will be described in subsection 7.6. In table 7-1, the file nn designation indicates the file code used to record the given report and is of no real concern to the user. In addition, a series of messages are produced which supply the user with information concerning special processing events that occurred during the execution of the data reduction program. Most of these processing messages are for information only, and can be ignored. The following subsections will describe all the reports listed in table 7-1, and subsection 7.5.25 will describe the processing messages that may be produced during the course of data reduction.

7.5.1 System Configuration and Channel Usage Report (File 42). This report documents the system identification, configuration, and the date and time of the monitoring period, as well as reporting the usage of all configured disk I/O channels. Tape Channel Usage is not reported by the Mass Store Monitor. Figure 7-1 is an example of this report. The heading line indicates the software version number that corresponds to this document. The version number should be 09-83 CHG-7. The first line after the heading provides the tape number(s) the report was generated from, the system identification, the date (in the form year, month, and day - YYMMDD), and the start and stop times (HH:MM:SS) of the MONITORING SESSION. The next several lines of output describe the overhead of all GMF

```
*************************************************************
* * *        M A S S   S T O R A G E   M O N I T O R    * * *
* * *             VERSION 09-83  CHG-7                  * * *
*************************************************************

        TAPE # 24502
   SYSID               DATE       START TIME    STOP TIME

OSCC2 ,6.4.ID        82-01-08      13:49:04      14:31:40       FOR A TOTAL OF 0.71 HOURS

              MONITOR            TIME(SEC)          % OVERHEAD
               EXEC                 100                 1.96
               MUM                  120                 2.35
               MSM                  160                 3.13
               NAME                  40                  .78
               FMS                   50                  .98
               TOTAL                                    8.42

CONFIGURATION:  DUAL    PROCESSOR 6680,  DUAL  IOM,  512K  MEMORY - 52 OF WHICH WERE HCM
                     THE HCM DOES NOT INCLUDE 3K FOR .CALC AND 6-8K FOR FILSYS
   PROCESSOR OPTIONS - 2K CACHE, EIS OPTIONS INSTALLED, FAST MEMORY ACCESS, INSTRUCTION OVERLAP ON
                    CPUS ACTUALLY CONFIGURED = 2.00  CPUS ACTUALLY AVAILABLE = 1.75
                    MAX # OF SYSTEM INTERCOM IOS ALLOWED ARE  128
                    MAX # OF OUTSTANDING SYSTEM INTERCOM IOS GENERATED WAS   26
                                     IOM NUMBER 0

 CHANNEL            TYPE          CROSSBAR       CONNECTS
  0-08             .DS450           0-09           47750
                                    1-08
                                    1-09
  0-09             .DS450         SEE ABOVE        16851
  0-12             .DS191           0-13           55524
                                    1-12
                                    1-13
  0-13             .DS191         SEE ABOVE         2479
  0-16             .DS181           0-17            8290
                                    1-16
                                    1-17

     (IN THE ACTUAL REPORT A SUMMARY OF IOM NUMBER 1 REPORT WOULD FOLLOW)
```

Figure 7-1. System Configuration and Channel Usage Report monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM".

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as:

$$\frac{\text{(CPU TIME Used by Monitor)}}{\text{(TOTAL Elapsed Time)x(Number of Processors)}}$$

Following the overhead description are six lines of configuration information describing the number of processors, IOMs, and amount of memory configured to the system. In addition, the size of GCOS Hard Core, the size of the Core Allocator and the size of FILSYS is also presented. The fourth line of the configuration data indicates the number of processors actually configured and actually available. These numbers might be different than shown on the first line due to the assigning and releasing of processors. In figure 7-1, we see that one processor was released for a period of time (i.e., CPUs actually available is equal to 1.75). The actual time that processors were available or released is indicated in the status message printouts (see subsection 7.5.25).

The final two lines of the report indicate how many simultaneous intercom I/Os are permitted and the maximum number of outstanding intercom I/Os recorded during the monitoring session. The intercom I/O feature is the means by which two programs residing in the H6000 can pass data to one another. If the system exhausts all available intercom I/O buffers, then the programs requiring this facility will be delayed. These two figures would provide an indication that the system did indeed exhaust all available intercom I/O buffer space whenever both lines contain the same number. Whenever this occurs, it shows that the entire buffer table was exhausted and therefore, the probability is high that jobs are being delayed while they are waiting for buffer space to free.

In addition, whenever the number of outstanding intercom I/Os is found to be equal to the total buffer capability, a warning message will be printed on the status message printout file. This message will indicate the time of day that the buffer pool was exhausted and a second printout will occur whenever the number of outstanding intercom I/Os falls back below the maximum allowed.

The next portion of the report documents the channel configuration by IOM, listing each configured disk channel number, the disk device type configured to that channel, and the channel crossbarring. The crossbar column shows those channels that are crossbarred to the channel identified under the channel column. If SEE ABOVE is found, the crossbarring has been displayed on a preceding channel. The I-CC format displayed under the CHANNEL heading identifies the IOM and the channel number. The last column of this report displays the number of all connect types issued over that channel. This section will be repeated for each IOM configured to the system. Figure 7-1 only displays IOM 0 activity. This report is always generated and cannot be turned off.

THIS PAGE LEFT INTENTIONALLY BLANK

```
***************************************************************
* * *       M S M   S Y S T E M   S U M M A R Y   R E P O R T       * * *
***************************************************************

TOTAL CONNECTS TO DSS181          8540 OF    192774 ( 4%) AT A RATE OF     12024. PER HOUR
TOTAL CONNECTS TO DSS191         76949 OF    192774 (40%) AT A RATE OF    108341. PER HOUR
TOTAL CONNECTS TO MS0450        107273 OF    192774 (55%) AT A RATE OF    151036. PER HOUR
TOTAL CONNECTS TO MPC-CONTROL       12 OF    192774 ( 0%) AT A LATE OF        16. PER HOUR
TOTAL CONNECTS TO IOM-0         130894 OF    192774 (68%) AT A RATE OF    184357. PER HOUR
TOTAL CONNECTS TO IOM-1          61880 OF    192774 (32%) AT A RATE OF     87154. PER HOUR

                         COMMANDS PER CHANNEL

                     COMMAND              COUNT
                  01 CONTRL                  12
                  17 FMT TK               15485
                  25 READ                126899
                  26 RD CR                  192
                  31 WRITE                35483
                  33 WR-VER               12862
                  40 RST ST                1841
                  TOTAL                  192774
```

Figure 7-2. MSM System Summary Report

7.5.2 System Summary Report (File 42). The System Configuration and Channel Usage Report and the System Summary Report may be used to assess overall system utilization. Figure 7-2 is an example of the System Summary Report. The first set of lines shows the number of connects to each monitored mass storage subsystem compared to the total connects issued to all Mass Storage subsystems and the connect rate per hour over each subsystem. Most systems will show a small number of Control Connects being generated by the MPCs configured to the system. These Control Connects will be summed together and listed as a separate subsystem line. Analysis on a Shared Mass Storage System shows the numoer of MPC connects generated to be a significant percentage of the total connects generated. The next lines show the breakdown of the mass storage connects by the IOM over which they were issued. The final part of this report is a list of the commands (octal code and mnemonic) issued to the mass storage subsystem and the count of each issued during the monitoring session. This report is always generated and cannot be turned off.

A well performing system, under a heavy workload, should show a high utilization of the configured resources. Figure 7-2 shows that the I/O activity is predominantly on the MSU450 subsystem configured on channels 8 and 9 of IOM 0 and 1 (see figure 7-1). The MSU450s are receiving 55% of all connects and, therefore, should be the major area of concern. The access rate for every subsystem is reported on the top of the System Summary Report and it can be seen that the MSU450s have an access rate significantly higher than the other subsystems. All signs indicate that if system throughput is being affected by disk activity, then the MSU450s would be the probable cause of such problems.

The next item to check should be the channel usage. The two highest used logical channels of any subsystem should be on a separate PSI channel of a two-PSI channel subsystem. If the highest used logical channels are not on separate PSI channels, the $ XBAR card in the startup configuration section is suspected as the cause. The channels are used in the order given on the $ XBAR card (i.e., if the primary channel is busy, the next channel tried is given on the crossbar). The alternate use of PSI channels for maximum simultaneity must, therefore, be appropriately specified in the boot deck. Subsection 8.3 provides a detailed explanation for analyzing the correctness of the crossbar configuration.

While looking at the System Summary Report, it is also of interest to note the ratio of READ commands to WRITE commands (over three to one in this example). This gives an indication of the nature of the usage of the mass storage space. A quick look at the number of write/verify (WR-VER) commands executed is also of interest as they are essentially double (WRITE, then READ) data transfer commands which require more device and channel time.

The general fraction of utilization for each logical channel gives an indication of the degree of simultaneity of access to the subsystem. If only N of the configured logical channels have nonzero counts, then there were never more than N accesses being performed simultaneously by the subsystem. The proportional relationships among the counts of accesses made over each of the logical channels are quantitative indications of the frequency of occurrence of specific levels of simultaneity. As an example, assume that the IOM-01 data reported in figure 7-1 (not shown in figure) indicated that 4487 connects out a total of 107,273 connects went to Channel 9, IOM 1. This means that only 4487 times during the measuring period were all four MSU450 disk channels being utilized simultaneously. In this example, channel queuing (i.e., shortage of channel power) would not appear to be a problem. This is not to infer that device queuing is not a problem, just that channel queuing does not appear to be a problem. If the number of accesses to the lowest priority channel is a larger percentage of the total accesses, then channel queuing needs to be examined. Queuing for devices and/or channels can be analyzed by running the Channel Monitor Data Reduction Program (see section 8).

7.5.3 <u>System Traces Captured by Monitor Report (File 42)</u>. This report contains the number of occurrences of each specific trace type recorded on the data collector tape processed by the MSMDRP (figure 7-3). This report provides little, if any, information required by the user for his analysis. This report is always generated and cannot be turned off.

7.5.4 <u>Channel Status Changes Report (File 29)</u>. This report lists the initial status for all tape and disk channels configured to the system (figure 7-4). If, during the course of the monitoring session, a given channel or IOM was dropped or added to the system (dynamic reconfiguration) a new report will be produced indicating the activation or deactivation changes and the time that the change occurred. Finally, this report will indicate whether the SSA cache option and FMS cache option are active, and if so, will indicate their initial status and any changes that occur to that status. If a given option is not active, a zero will be reported for each of the values. This report is always generated and cannot be turned off.

7.5.5 <u>Physical Device, Device ID Correlation Table (File 42)</u>. Each mass storage device configured in the system is listed with a unique device ID. A typical report is presented in figure 7-5. This unique device is needed since different devices can have the same device number on the Honeywell 6000. (See Device ID 1, Device ID 7, and Device ID 18 of figure 7-5). These unique numbers are referenced in several reports produced by the MSMDRP. This report is always generated and cannot be turned off.

7.5.6 <u>Device Space Utilization Report (File 42)</u>. The device space utilization histogram report is produced for every device on the mass storage subsystem and shows the distribution of access to the device space. Figure 7-6 is an example. It should be noted that the name of the device is also given. This example presents all connects made to the device with the name RF5. If an exchange took place and the RF5 disk pack was moved from 0-08-05 to 0-08-01 the data reduction program will account for that exchange and any connects that are made to 0-08-01 will be reported on this histogram and not to the 0-08-01 histogram. Entries in the column headed CYLNDR NUMBER give the range of cylinders which form each histogram bucket. The number of cylinders in each bucket is a function of the device type. The entries in the column headed INDIV. NUMBER give the number of accesses made to that device within the physical space defined by the range of cylinders.

Similarly, the columns headed INDIV. PRC and CUMUL. PRC give the individual and cumulative percentages of all accesses to that device which were made within each cylinder range. The graphic portion of the display gives a visual indication of the percentage of accesses which were made for each range of the device space. This helps to quickly assess the access pattern of the usage for the device, i.e., whether the device is totally allocated and used or locally used. Figure 7-6 shows a device whose usage is split between two extremes.

In the upper right hand corner of the report, a report number is indicated. This report number is used only to distinguish one histogram from another and in no way indicates the device to which the report refers. In addition, report numbers may not appear sequentially and this, in no way, is indicative of a problem. This report is always generated and cannot be turned off.

7.5.7 Device Seek Movement Report (File 42). The seek movement histogram is produced for devices in the mass storage subsystem being analyzed and provides the distribution of distance traveled by the read mechanism. Figure 7-7 is an example. The data used to generate this report is the absolute value of the difference between the cylinder addresses of each successive access to the given device. The column headed CYLNDR MOVED contains the range of seek movement distance for each line of the report. The column headed INDIV. NUMBER contains the counts of the number of accesses which caused the arm to be moved that distance. Figure 7-7 shows 854 accesses caused no arm movement (the same cylinder was successively accessed) for IOM-0 Device 05 on PUB 8. The INDV. PRC and CUMUL. PRC columns give the individual and cumulative percentages of the accesses to that device which resulted in a particular range of seek movement.

When using this report to analyze seek elongation problems, it is fairly easy to determine the location of files that are causing the seek elongation problem. For example, figure 7-7 shows 263 (12.8 percent) of the accesses to the device resulted in an arm movement of between 714-730 cylinders. If figure 7-6 is examined, it will be observed that 564 connects were made to cylinder 0 and 213 connects were made to cylinders 714-730. If the accesses to these 2 areas were made in alternating fashion, the resulting arm movement distance would be 714-730 cylinders. In addition, figure 7-6 also shows 207 accesses to cylinders 34-50 and 612 accesses to cylinders 748-764. Once again, if accesses to these two areas were made in alternating fashion, the resulting arm movement distance would be 714-730 cylinders. By using other reports produced by MSMDRP, the analyst can next determine the actual files located in these areas, and perhaps relocate one or more files so as to eliminate this excessive seeking.

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-05-- MS0450 RF5

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC.	INDIV. PRC.	CYLNDR NUMBER	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	REPORT 50
854	854	41.801	41.801	0- 0	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx	
459	1313	64.268	22.467	1- 16	Ixxxxxxxxxxxxxxxxxxxxxx	
118	1431	70.044	5.776	17- 33	Ixxxxxx	
63	1494	73.128	3.084	34- 50	Ixxx	
17	1511	73.960	0.832	51- 67	Ix	
5	1516	74.205	0.245	68- 84	I	
7	1523	74.547	0.343	85- 101	I	
10	1533	75.037	0.489	102- 118	I	
1	1534	75.086	0.049	119- 135	I	
18	1552	75.967	0.881	136- 152	Ix	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
10	1633	79.931	0.489	663- 679	I	
36	1669	81.694	1.762	680- 696	Ixx	
2	1671	81.791	0.098	697- 713	I	
263	1934	94.665	12.873	714- 730	Ixxxxxxxxxxxxx	
2	1936	94.763	0.098	731- 747	I	
107	2043	100.000	5.237	748- 764	Ixxxxx	

2043 ENTRIES TOTAL AVERAGE = 167.63534 VARIANCE = 85389.356 STANDARD DEVIATION = 292.215

Figure 7-7. Device Seek Movement Report

The above procedure becomes more complicated when analyzing disk types 500 and 501. This device type consists of two logical devices configured as one physical device. Access to one member of the pair may affect the performance of the other member because of the fact that the hardware design considers the two devices to be a single physical device. In figure 7-7.3, it can be observed that 9 accesses to that device resulted in an arm movement of 545-561 cylinders. If figure 7-7.2 is examined next, it can readily be seen that the access pattern to this device could never result in an arm movement of between 545-561 cylinders. However, if figure 7-7.1 is checked, it will be observed that 7918 connects did occur on cylinders 545-561. Therefore, we now see an example of seek movement conflict between device pairs, as opposed to conflict on a single device. In order to analyze this situation, it would be necessary to determine the files located at cylinders 545-561 on device MSH and also the files located at cylinder 0 on device MSI.

Further confirmation of a seek elongation problem could be found by analyzing the Head Movement Efficiency Report (see subsection 7.5.8 and figure 7-8). If a problem exists then the connects/arm movement column for a particular device should approach a value of one. Figure 7-8 indicates how many connects are issued between each movement of the arm. Therefore, even though we may have long seeks occurring on the device, if a large number of connects are being processed between these seeks, this would tend to lessen the impact of the long seeks.

For the Device Space Utilization Report and Device Seek Movement Report, an entry is made only for multi-command connects (see subsection 7.3) such as a seek/read or seek/write. If the first command of an IO connect is not a seek, or pre-seek, then an entry will not be made to this set of histograms. For this reason, the number of connects reported in these reports, for a given device, may be somewhat lower than that reported in the Proportionate Device Utilization Histogram, described in subsection 7.5.20.

7.5.8 <u>Head Movement Efficiency Report (File 42)</u>. This report displays how many connects are issued per arm movement of the device. Figure 7-8 (generated from a different monitoring session than that used to obtain figures 7-6, 7-7, 7-7.1, 7-7.2 and 7-7.3) is an example. The first three columns give the IOM, Channel, and Device number of the device. This is followed by the number of connects issued to that device and the number of times any arm movement was required (size of the seek is not considered). The final column indicates the ratio of connects to arm movements. The larger this ratio is, the more efficient is the device (i.e., the larger is the number of connects being handled between each arm movement of the device). Following the breakdown of arm movement by individual device, a summary is presented for arm movement within each individual mass storage subsystem. This is followed by three lines of output summarizing the overall efficiency of the entire disk subsystem. The first line presents the total number of connects issued, the second

```
                    DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                         SPACE UTILIZATION OF IOM-0,PUB-16,DEVICE-09-- MS0500          MSH

INDIV.   CUMUL.   CUMUL.   INDIV.   CYLNDR                PERCENT OF OCCURRENCE                REPORT
NUMBER   NUMBER   PRC.     PRC.     NUMBER    00  10  20  30  40  50  60  70  80  90  100    86
                                              I....I....I....I....I....I....I....I....I....I....I
  163      163     1.195    1.195    32-  34  Ix
  335      498     3.650    2.455    35-  51  Ix
  379      877     6.427    2.778    52-  68  Ix
   .        .        .        .       .    .  .
   .        .        .        .       .    .  .
   .        .        .        .       .    .  .
    0     4180    30.634    0.      511- 527  I
    0     4180    30.634    0.      528- 544  I
 7918    12098    88.663   58.029   545- 561  Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxx
   57    12155    89.080    0.418   562- 578  I
    0    12155    89.080    0.      579- 595  I
   .        .        .        .       .    .  .
   .        .        .        .       .    .  .
   .        .        .        .       .    .  .
  124    12279    89.989    0.909   732- 748  I
    4    12283    90.018    0.029   749- 765  I
   88    12371    90.663    0.645   766- 782  I
  575    12946    94.877    4.214   783- 799  Ixx
  699    13645   100.000    5.123   800- 816  Ixxx

13645 ENTRIES TOTAL    AVERAGE =  450.71147   VARIANCE =  53246.213   STANDARD DEVIATION =  230.751
```

Figure 7-7.1. Space Utilization Report for Device 9

```
                    DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                         SPACE UTILIZATION OF IOM-0,PUB-16,DEVICE-10-- MS0500          MSI

INDIV.    CUMUL.    CUMUL.    INDIV.    CYLNDR              PERCENT OF OCCURRENCE              REPORT
NUMBER    NUMBER    PRC.      PRC.      NUMBER    00   10   20   30   40   50   60   70   80   90   100   88
                                                   I....I....I....I....I....I....I....I....I....I....I
  20        20      60.606    60.606     0-   0   Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
   0        20      60.606    0.         1-  17   I
   0        20      60.606    0.        18-  34   I
   .         .        .        .          .    .   .
   .         .        .        .          .    .   .
   .         .        .        .          .    .   .
   0        20      60.606    0.       494- 510   I
   5        25      75.758    15.152   518- 527   Ixxxxxxxx
   0        25      75.758    0.       528- 544   I
   .         .        .        .          .    .   .
   .         .        .        .          .    .   .
   .         .        .        .          .    .   .
   0        25      75.758    0.       698- 714   I
   5        30      90.909    15.152   715- 731   Ixxxxxxxx
   3        33     100.000     9.091   732- 748   Ixxxxx

 33 ENTRIES TOTAL    AVERAGE =   254.39394    VARIANCE =   103593.147    STANDARD DEVIATION =   321.859
```

Figure 7-7.2. Space Utilization Report for Device 10

```
                    DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                         SEEK MOVEMENT OF IOM-0,PUB-16,DEVICE-10-- MS0500          MSI

INDIV.    CUMUL.    CUMUL.     INDIV.    CYLNDR                  PERCENT OF OCCURRENCE                REPORT
NUMBER    NUMBER    PRC.       PRC.      NUMBER    00   05   10   15   20   25   30   35   40   45   50    89
                                                   I....I....I....I....I....I....I....I....I....I....I
  14        14      42.424     42.424      0-   0  Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
   2        16      48.485      6.061     13-  17  Ixxxxxx
   0        16      48.485      0.        18-  34  I
   0        16      48.485      0.        35-  51  I
   0        16      48.485      0.        52-  68  I
   1        17      51.515      3.030     69-  85  Ixxx
   .         .         .          .        .    .  .
   .         .         .          .        .    .  .
   .         .         .          .        .    .  .
   0        21      63.636      0.       511- 527  I
   0        21      63.636      0.       528- 544  I
   9        30      90.909     27.273    545- 561  Ixxxxxxxxxxxxxxxxxxxxxxxxxxx
   0        30      90.909      0.       562- 578  I
   0        30      90.909      0.       579- 595  I
   0        30      90.909      0.       596- 612  I
   .         .         .          .        .    .  .
   .         .         .          .        .    .  .
   .         .         .          .        .    .  .
   2        32      96.970      6.061    715- 731  Ixxxxxx
   0        32      96.970      0.       732- 748  I
   0        32      96.970      0.       749- 765  I
   0        32      96.970      0.       766- 782  I
   0        32      96.970      0.       783- 799  I
   1        33     100.000      3.030    800- 816  Ixxx

 33 ENTRIES TOTAL   AVERAGE =  247.93939    VARIANCE =  80552.542    STANDARD DEVIATION =  283.818
```

Figure 7-7.3. Seek Movement Report for Device 10

THIS PAGE LEFT INTENTIONALLY BLANK

MSM HEAD MOVEMENT EFFICIENCY REPORT FOR SYSTEM HMCC ON 81-12-07

IOM	PUB	DEVICE	CONNECTS	ARM MOVES	CONNECTS/ARM MOVEMENT
0	08	01	34848	28491	1.223
0	08	02	21479	8673	2.477
0	08	03	30030	18410	1.631
0	08	04	31128	19449	1.600
0	08	05	0	1	0.
0	08	06	220	1	220.000
0	12	01	6449	3138	2.055
0	12	02	2520	86	29.302
0	12	03	2048	549	3.730
0	12	04	3303	562	5.877
0	12	05	4818	1887	2.553
0	12	06	2929	754	3.885
0	12	07	220	1	220.000
0	12	08	4824	1479	3.262
0	12	09	9145	3303	2.769
0	12	10	3249	690	4.709
0	12	11	2285	754	3.031
0	16	01	0	1	0.
0	16	02	0	1	0.
0	16	03	12	1	12.000
0	16	04	3754	168	22.345
0	16	05	260	12	12.667
0	16	06	220	1	220.000
0	16	07	220	1	220.000

DEVICE TYPE	CONNECTS	ARM MOVES	CONNECTS/ARM MOVEMENT
.DS181	4466	185	24.141
.DS191	41790	13203	3.165
.DS450	117705	75025	1.569

TOTAL MASS STORAGE CONNECTS	163961
TOTAL ARM MOVEMENTS	88413
HEAD MOVEMENT EFFICIENCY	1.854

Figure 7-8. Head Movement Efficiency Report

line presents the total number of arm movements, irrespective of the number of cylinders traversed by each movement, and the final line presents the overall head movement efficiency (line 1 divided by line 2). By examining figure 7-8, the user will observe that the MSU450s appear to be least efficient, with the exception of device number 2. This report is always generated and cannot be turned off.

7.5.9 System File Use Summary Report (File 21). This report indicates where each system file is located and to what extent it was accessed across the measurement session. Figure 7-9 shows an example of this report. This report is produced by default but can be turned off with the input option (OFF) (subsection 7.6.9). The sum of accesses to all system files is then expressed as a percentage of all mass storage accesses. The system files are those files defined, via startup in the .CRDIT table. As can be seen in figure 7-9, the file names listed under the File Name column are not the actual file name, but rather relative file names indicative of their position description within the startup deck. Actual file names can be output in this report if the user selects the input option described in subsection 7.6.4.

Included in the list of system files are two additional file names of ACTNG and SYSOUT. These files indicate accesses to the accounting file and any configured SYSOUT files. For these files, the information provided under STARTING SECTOR/CYLINDER is not the actual starting address, but rather the smallest address accessed during the monitoring session. Likewise, the information provided under the LENGTH column is not the actual length of the file, but rather the difference between the largest and smallest addresses accessed during the monitoring session.

This list of system files is then followed by a list of modules which reside in hard core because they are hard core modules or because they have been loaded into hard core by system personnel in order to save on I/O processing. If a system module is not loaded in hard core, and is required for some processing, then the system must perform an IO function to read this module from disk into a user's SSA work space. A significant amount of such system IO can cause severe system degradation. This degradation can be reduced by placing additional system modules in hard core or else by increasing the size of SSA Cache Memory. The System File Use Summary Report and the following Individual Module Activity Report should provide sufficient information to determine whether user action is required to reduce system IO overhead. If the percentage of system IO reported in this report is greater than 5-7%, then some user action is probably required. If addtional hard core space is available, the user should move as many system modules as possible into hard core. Each hard core module requires 1/2K (512 words) of memory, and there is 64K of hard core memory available. The System Configuration and Channel Usage Report (subsection 7.5.1) indicates the amount of hard core currently in use. The Individual Module Activity Report (subsection 7.5.10) can be used to indicate which system modules should be transferred to hard core. If sufficient hard core is not available, then the size of SSA cache should probably be increased. Once again, the Individual Module Activity Report should be referenced to see if this type of action would aid in reducing system IO.

MSM SYSTEM FILE USE SUMMARY REPORT FOR SYSTEM NMCC ON 83-08-02

FILE NUMBER	FILE NAME	IOM-PUB-DEV	STARTING SECTOR/CYLINDER		LENGTH (SECTOR)	ACCESSES
1	SYSTEM FILE1	0- 8- 1	12260/	16	5000	0
2	SYSTEM FILE2	0- 8- 5	60/	0	7500	680
3	SYSTEM FILE3	0- 8- 1	17260/	22	1250	50
4	SYSTEM FILE4	0- 8- 3	12820/	16	15000	0
5	SYSTEM FILE5	0- 8- 2	6260/	8	8500	1621
6	SYSTEM FILE6	0- 8- 3	1820/	2	11000	214
7	SYSTEM FILE7	0-12- 8	60/	0	11000	0
8	SYSTEM FILE8	0- 8- 5	7560/	9	11000	1361
9	SYSTEM FILE9	0- 8- 4	60/	0	5500	0
10	SYSTM FILE10	0- 8- 4	5560/	7	21000	359
11	SYSTM FILE11	0- 8- 1	18520/	24	1000	0
12	SYSTM FILE12	0- 8- 5	18560/	24	12500	0
13	SYSTM FILE13	0- 8- 4	26560/	34	7500	0
14	SYSOUT	0-13-12	9275/	12	188505	64
15	SYSOUT	0-13- 4	42620/	56	5	4
16	ACTNG	0- 8- 5	34405/	45	675	136
17	SYSOUT	0- 8-21	430590/	566	1717	27
18	SYSOUT	1-12- 1	2199/	2	25681	199
19	SYSOUT	0-12-10	7880/	10	5	472
					TOTAL	5187(8%)

```
.MBRT1   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MCPIO   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MDISP   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MDNET   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MDUMP   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MFALT   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.MGEPR   LOADED BY STARTUP   TYPE IS   HRD CORE MOD
.
.
.
```

Figure 7-9. System File Use Summary Report

7.5.10 Individual Module Activity Report (File 21). This report shows the accessing done to each system module (figure 7-10). The report presents the system file the module resides in, followed by the module name and type. The module location, access count, and percentage of system file usage is then given. The last two entries give the total number of SSA CACHE buffer hits and disk loads this module accumulated (these values are 0 if SSA CACHE is not active). A minus one in column indicates a nonstandard SSA module. The Number of Accesses column reports the number of connects made to this SSA module as determined by the issuance of a trace type 7. The Disk Load column reports the number of times the SSA CACHE logic claims to have issued a connect to this SSA module. In those cases where no lost data occurred during the monitoring session and a GMC termination record was generated (i.e., GMC terminated correctly), these two columns should display equal values. Figure 7-10 shows this to be the case for almost all of the SSA modules. However, there are some exceptions. Module .MALC6 shows 109 accesses but only 104 disk loads (i.e., a difference of five). This apparent inconsistency has been reported to Honeywell, and an explanation requested.

If lost data occurred during the monitoring session (i.e., trace type 7 data lost), the Number of Accesses column could be significantly lower than the Disk Load column. If, on the other hand, GMC did not abort cleanly and a termination record was not generated, the Disk Load column could be significantly lower than the Number of Accesses column.

The module names are all hard coded in the MSM Data Reduction Program. To verify that the module names are correct, especially for commercial systems, the following GMAP program should be run to produce a module list and number. This list can then be compared to the list in the routine BLOCK DATA. The GMAP program is as follows:

```
$     GMAP   NDECK
$     LODM   .G3MAC
$     LODM   .G3MCR
$     .MDEF1,LIST
$     SYMDEF START
START NULL
$     END
```

Finally, it should be noted that data from the last two columns pertains only to "STANDARD SSA" modules (see the Type column). Modules that are typed as "ABSOLUTE" or "EXCEPT PROC" are not placed into SSA Cache Core and therefore do not generate values for the last two columns. This report is produced by default but can be turned off with the input option (OFF) (subsection 7.6.9).

When generating this report, the data reduction program writes to file code 55, which is used to produce a Job SSA Module Usage Report (subsection 7.5.11). If this report is desired, the report must be requested via input option (MODULE) (subsection 7.6.6).

AD-A138 533 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 7 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 01 OCT 83 CCTC-CSM-UM-246-82-CHG-7

2/3

UNCLASSIFIED F/G 9/2 NL

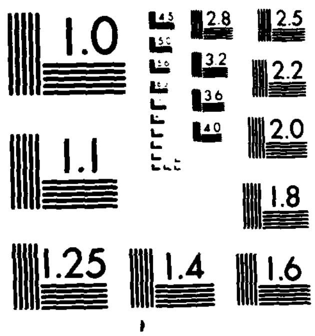

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

At the bottom of this report, a summary line is produced indicating the percentage of buffer hits and disk loads. If the percentage of buffer hits is less than 90, then the size of SSA cache should be increased. For each 1K increase, 2 additional modules will be loaded into the SSA Cache memory.

When using this report to determine which additional SSA modules should be placed into GCOS Hard Core, the user should reference the "% of Activity" column. Those modules with the largest reported figure would be candidates for movement. In figure 7-10, .MALC6, .MALC9, .MFS03, .MFS04 would be candidates for movement.

THIS PAGE LEFT INTENTIONALLY BLANK

MSM INDIVIDUAL MODULE ACTIVITY REPORT FOR NMCC ON 81-12-07

SYSTEM FILE	MODULE NAME	TYPE	IOM-PUB-DEVICE	SECTOR IN FILE	NUMBER ACCESSES	% OF ACTIVITY	BUFFER HITS	DISK LOADS
SYSTEM FILE4	.MAC02	STANDARD SSA	0- 8- 1	180	44	1	42	44
SYSTEM FILE4	.MALC2	STANDARD SSA	0- 8- 1	452	11	0	23	11
SYSTEM FILE4	.MALC5	STANDARD SSA	0- 8- 1	473	21	0	18	21
SYSTEM FILE4	.MALC6	STANDARD SSA	0- 8- 1	482	109	2	77	104
SYSTEM FILE4	.MALC9	STANDARD SSA	0- 8- 1	501	101	2	83	95
SYSTEM FILE4	.MBRT2	STANDARD SSA	0- 8- 1	670	12	0	3	12
SYSTEM FILE4	.MBRT3	STANDARD SSA	0- 8- 1	679	3	0	0	0
SYSTEM FILE4	.MBRT5	ABSOLUTE PRG	0- 8- 1	697	2	0	0	0
SYSTEM FILE4	.MBRT6	ABSOLUTE PRG	0- 8- 1	729	26	0	0	0
SYSTEM FILE4	.MCAL1	STANDARD SSA	0- 8- 1	831	20	0	858	20
SYSTEM FILE4	.MCAL2	STANDARD SSA	0- 8- 1	840	1	0	0	1
SYSTEM FILE4	.MFLT1	STANDARD SSA	0- 8- 1	1332	78	1	1402	77
SYSTEM FILE4	.MFS03	STANDARD SSA	0- 8- 1	1361	101	2	182	93
SYSTEM FILE4	.MFS04	STANDARD SSA	0- 8- 1	1370	99	2	30	93
SYSTEM FILE4	.MFS06	STANDARD SSA	0- 8- 1	1388	19	0	4	19
SYSTEM FILE4	.MFS08	STANDARD SSA	0- 8- 1	1406	18	0	37	18
SYSTEM FILE4	.MFS09	STANDARD SSA	0- 8- 1	1415	57	1	526	52
• • • •								
				TOTAL			76357	33455
				% HITS, LOADS/HITS + LOADS			69.53	30.47

Figure 7-10. Individual Module Activity Report

7.5.11 SSA Module Usage Report by Job (File 21). When requested by the user, this report will produced a listing for every job run during the monitoring period, showing all SSA modules referenced by that job and the number of such references. An example is shown in figure 7-11. This is the best method to use when determining which SSA modules should be softloaded into TSS core. This also provides an excellent means for studying the usage of SSA modules in general. This report is off by default and must be requested with a user input option (MODULE) (subsection 7.6.6).

7.5.12 File Code Summary Report (File 23) (NAME=FILECODE). The File Code Summary Report lists, by each activity, the files allocated to mass storage, their location and size, and the number of accesses made to each in the system during the monitoring period. Figure 7-12 is an example of this report. The activities in this report are in the same order as they appear in the Activity Summary Report (see subsection 7.5.15).

Each activity is identified by its SNUMB, activity number, and $ IDENT and USERID cards. There are as many data lines as necessary to describe each mass storage file used by the activity and the number of times the file was accessed. There is one line per file, and the file is described by its two-character file code, the device on which it was allocated (ALLOCATED DEVICE), its origin on that device (FILE ORIGIN) in units of LLINKS (320 words) and cylinders relative to the beginning of the device, and the size of the file (FILE SIZE) in LLINKS and cylinders. The column headed CONNECTS gives the count of the number of accesses made to the file.

There are several special file codes that will appear in this report. The following file codes will appear for almost every activity that is processed:

00 - Mass storage accesses made without a normal PAT entry; e.g., accesses made by the operating system as part of job intialization which are done without a PAT for efficiency

-- - All accesses to SYSOUT

=1 - File grow, PAT refresh, permanent allocation

=7 - Temporary allocation

=9 - SWAP/GECALL

0% - Load, pop, or push of an SSA

*, - FILSYS catalog search connect

J* - JCL file

*J - Data subfiles for a job

=A - Accesses to the accounting file - normally made by $CALC

DATA COLLECTED ON 81-12-25 AT 18:06

SNUMB	MODULE	# ACCESSES	MODULE	# ACCESSES	MODULE	# ACCESSES	MODULE	# ACCESSES
3625T	.MCAL1	7	DIRXLT	1	DDLXLT	1	.MFLT1	15
3625T	.MFLT2	11	.MMORE	5	.MAC10	7	.MFS09	17
3625T	.MFS03	8	.MFS36	4	.MFS15	4	.MFS61	8
3625T	.MFS73	10	.MFS37	4	.MRELS	4	.MREL1	3
3625T	.MFS04	24	.MFS13	2	.MSY11	11	.MSYT2	2
3625T	.MSYT3	2	.MIOS1	3	.MBRT2	3	.MFS27	2
3625T	.MGP23	2	.MBRT6	2	.MALC9	3	DBUG	1
3625T	.MMOR1	1	.MFS19	1	.MFS34	1	.MALC5	1
3625T	.MFS49	1	.MALC7	1	.MFS23	1	.MFS18	1
3625T	.MSNP1	2						
3637T	.MCAL1	6	ECFR	1	.MFS09	4	.MFLT1	2
3637T	.MFLT2	1	.MFS08	6	.MFS68	2	.MSYT2	5
3637T	.MSYT3	2	.MIOS1	1	.MBRT2	1	.MSYT1	4
3637T	.MFS27	1	.MGP23	1	.MBRT6	1	.MALO9	2

Figure 7-11. SSA Module Usage Report By Job

FILE CODE SUMMARY REPORT FOR SYSTEM NMCC ON 81-12-01

SNUMB,ACTVY #,IDENT,USERID 2802T- 1 1820513/10/4804/FORDI,KONKEL 674IDP00

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)
00	479	0/ 0	1-10- 2	0/ 0
*1 T S N F	76	48/ 1	0-12- 1	44361/291
B* T S N F	145	72/ 1	0-12- 9	26085/171
-- P R N F	53	0/ 0	0-12- 7	0/ 0
*Z T R N F	83	60/ 1	0-12- 4	48778/320
*C P S C F	44	139/ 1	1- 8-13	23523/154
S* T S N F	157	72/ 1	0-12- 4	48263/317
K* T S N F	36	600/ 4	2- 9-10	50649/333
D* P S N F	8	1/ 1	1- 8- 9	46522/306
*5 T S N F	38	36/ 1	0-12-10	29611/194
R* P S C F	61	65/ 1	0-13- 8	23364/153
*4 T S N F	2	12/ 1	0-12-11	48662/320
*J T R N F	1	0/ 0	1-10- 9	46518/306

SNUMB,ACTVY #,IDENT,USERID 2760T- 3 1820615/30/4955/GFEH* 220IDP11

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)
20 T R N F	1	420/ 3	0- 8- 4	51921/341
00	81	0/ 0	1- 8- 2	0/ 0
02 T R N F	384	180/ 2	0- 8- 8	61573/405
05 P S C F	3	14/ 1	0- 8- 7	25274/166
03 T S N F	20	1200/ 8	0- 8- 7	49663/326
01 T R N F	36	240/ 2	0- 8- 7	47671/313
-- P R N F	44	0/ 0	0- 8-12	0/ 0
10 T R N F	290	240/ 2	1-10-10	22349/147
04 P S N F	1	1/ 1	0- 8- 7	49383/324
12 T R N F	2	240/ 2	0- 8- 4	51681/340
09 T S N F	2	1200/ 8	0- 8- 8	46470/305
18 T S N F	2	60/ 1	1- 8-13	49109/323
15 T S N F	2	60/ 1	0-10-10	25549/168
19 T R N F	1	300/ 2	0- 8- 2	48910/321
*J T R N F	1	0/ 0	1- 8- 7	49378/324

Figure 7-12. File Code Summary Report

The file code report for the TSS subsystem will also contain several special file codes. Following is a list of these:

TU - user files requested under TSS subsystems.
   - file code referenced by TSS that contains a character that cannot be printed. The   symbol is being used instead.
JO - TSS I/O to this file code is JOUT processing.

It should be realized that these file codes are not necessarily used for a unique file request, and therefore may actually reference several different files. Therefore, in figure 7-12 it can be seen that activity 1 for SNUMB 2802T made 479 connects to file code 00 located on device 1-10-2. In actuality, the activity may have connected to several different devices, each time referencing this special file code. Instead of reporting these connects as separate entries, all the requests are grouped, and the allocated device that is reported is the the one that was referenced by the first of the 479 connects. There is a user option available (AREA) that expands the usage definition of these file codes. This option will then display all unique requests using these file codes rather than grouping all of the unique file codes into a single collection file code. In addition, it should be noted that the information reported under the file size column (for these special system files and TSS files) is normally reported as 0, and the information reported under the origin column (for these same files) is not the actual origin of the file, but rather the location of the first seek address made to that file.

This report may produce excessive output and therefore is not produced under default conditions. The user must explicitly request this report with the use of an input option (ON) (see subsection 7.6.10).

Each file code entry has information which indicates the type of file by a "P" for permanent and "T" for temporary. This entry is located on the right side of the file code characters. Immediately to the right of the file type character is the access characteristics of the file. This is denoted by an "R" if random, or an "S" for a sequential file. The next character defines a permanent file as cataloged, "C", or noncataloged, "N". This character for a temporary file will be a blank as it is neither cataloged nor noncataloged. GCOS files are usually defined as noncataloged, permanent files and are treated at startup time much like permanent files by the file system. They are not, however, given the allocation treatment during normal operation time that is given to permanent files. The final character is an "F" for fixed or an "R" for removable.

NOTE: A file size of 16384 for a nonsystem file implies the file size is greater than 16383 llinks.

7.5.13 Cat/File String Report (File 23). Immediately following the File Code Summary Report, the user will find the CAT/File String Report (figure 7-13). This report will provide the CAT/File String for every permanent user file referenced during execution. It will not list the CAT/File String for any special system files. In addition, it will indicate the total number of connects required to locate the file (catalog searching) every time the system was required to search for that file via a MME GEFSYE. Finally, it will indicate the number of connects required to

ACTIVITY SUMMARY REPORT FOR SYSTEM NMCC ON 82-06-25

JOB	MAX IO QUEUES	MAX IO IN PROCESS	MAX INTERCOM OUTSTANDING	CP TIME(MS)	CONNECTS MASS	CONNECTS TOTAL	CP TIME PER CONNECT	
$CALC- 0	20	9	1	5157	440/	440	11.72/	11.72
$PALC- 1	5	1	0	3195	906/	1030	3.53/	3.10
$SYOT- 0	17	6	0	0	512/	512	0. /	0.
$RTIN- 0	40	1	0	5141	3/	3	1713.67/	1713.67
TS1 - 1	105	7	0	51245	3281/	3281	15.62/	15.62
$LOGN- 0	50	1	0	1485	13/	13	114.23/	114.23
NCP - 1	87	9	5	76381	1114/	1114	68.56/	68.56
TLNT - 1	8	1	0	0	1/	1	0. /	0.
FTS - 1	183	20	18	83149	3151/	3151	26.39/	26.39
TLCF - 1	90	1	0	2	3/	3	0.67/	0.67
7088T- 1	5	2	0	2946	143/	151	20.60/	19.51
FSPMM- 2	8	1	0	0	2/	2	0. /	0.
MUM - 1	0	0	0	284	0/	47	0. /	6.04
SCHED- 0	0	0	0	10453	135/	149	77.43/	70.15
			TOTALS	239438	9704/	9897	24.67/	24.19

**** END OF REPORT SET NUMBER 1   1130-1141 *****

Figure 7-15. Activity Summary Report

This report will display all connects issued by a job with no regard to the type of I/O command or the validity of the seek address (see subsection 7.3).

7.5.16 Device Area File Code Reference Report (File 22). This report is generated to provide details on the jobs accessing a specific device area with their file codes. Figure 7-16 displays an example. The devices and areas to be listed are defined by the user when requesting input option Area (subsection 7.6.1). In figure 7-16, there are 10 areas requested for investigation. Each activity that accessed a device area is displayed in the report. At the end of the report, the number of connects found to each requested area is also given. This report is identical in format to the File Code Summary Report (subsection 7.5.12) except that this report contains only the file codes which referenced the specific area of the desired devices. The AREA N of each file code specifies within which area of the possible set of requested areas this particular file code fell. When this option is selected, the file code reference will automatically be expanded and special system file codes will be reported only if they actually referenced the requested area (see subsection 7.5.12). In addition, if the special system file codes referenced multiple areas, these file codes will appear multiple times within this report. In figure 7-16, it can be seen that activity 3 of job 52323 has multiple references to file code 0%. In the standard File Code Summary Report, all these references would have been grouped as a single reference, but in this report, they are expanded within each unique area requested by the user.

A complete explanation of the special file codes can be found in subsection 7.5.12. This report is not produced under default conditions and must be requested with a special user input option (AREA) (subsection 7.6.1).

7.5.17 Device File Use Summary Report (File 21). This report shows the device use by the accesses per file class (temporary or permanent). Figure 7-17 is an example of this report. Each of these classes of allocation is subdivided into sequential and random files and their corresponding percentage of the total file use is presented in the report. File 00 accesses are not included in this report. This report will reflect only multicommand connects. The device numbers being reported under the "DEVICE" column are the unique set of device numbers generated by the MSMDRP (see subsection 7.5.5). This report is on by default but may be turned off with a user input option (OFF) (subsection 7.6.9).

7.5.18 Chronological Device Utilization Report (File 26). This report provides a chronological listing of the six most active disk devices, by device number and their probability of utilization (see figure 7-18). This report is so designed that any time quantum can be set in the report. By varying the time quantum parameter, the user may select integer values from 1 to n (where n is a positive value in seconds). A time quantum variation is requested with a user input option (TIMEQ) (subsection 7.6.14).

followed by many accesses to device 3 followed by many accesses to device 2 followed by many accesses to device 1, etc. Each device could have been a bottleneck for a subperiod of the total monitoring period. This could also have been the case if the proportionate utilization of each device was equal. The Channel Monitor can be used to uncover this cyclic type of usage. In addition, the Chronological Device Utilization Report (see subsection 7.5.18) was designed to uncover this type of problem by breaking down device utilization over time, rather than by utilizing a histogram. Nevertheless, when a single or small number of devices has a disproportionately large share of the accesses, they are potential bottlenecks and their usage should be further analyzed.

This report will show all connects that were issued to a given device. This includes all read/write connects, as well as any command type connects issued to a given device.(See subsection 7.3).

This histogram can report a maximum of 50 devices. If a site is configured with more than 50 devices, a second report will be produced as a continuation of Report 1.

7.5.21 Elapsed Time Between Seeks Report (File 42). This is a histogram report for the frequency of occurrence of elapsed time intervals between the issuance of mass storage access connects. Figure 7-21 presents a sample. The elapsed time is calculated as the time difference between successive mass storage connects from the central system and thus is representative of the workload. It does not provide any meaningful information on the subsystem service capabilities.

The data presented give the count (INDIV. NUMBER) and percentage (INDIV. PRC) of elapsed time between accesses which fell within each time range. The column headed TIME MSECS gives the time range in milliseconds. Thus, the data of the row with a time of 18 gives the count and fraction of elapsed time intervals in the range of 17+ to 18 milliseconds. The columns headed CUMUL. NUMBER and CUMUL. PRC. give the accumulated counts and percentage and are useful in describing the mass storage rates, e.g., 75.4 percent of the accesses occur less than 21 ms after the last access.

The bottom of the report provides a statistical summary of the data in the report. Statistics given include average, variance, and standard deviation. These statistics apply to all data points that were measured. The statistics concerning OUT OF RANGE are for those data points which fall outside the range of the histogram. OUT OF RANGE points are included in the previous statistics. This report is always generated and cannot be turned off.

7.5.22 Data Transfer Size Report (File 42). A sample histogram report on the frequency of occurrence of sizes of the data blocks transferred between mass storage and main memory is given in figure 7-22. Refer to subsection

7.5.12 for a description of the histogram format. This report has increments of 64 words, and the number in the column headed NUMBER WORDS is the upper value. The occurrence of certain data transfer sizes should be anticipated. For example, 64-word blocks are used for catalog accessing; in other parts of GCOS, standard system format is 320 words. SSA modules are usually slightly less than 512 words. When the Timesharing Subsystem

THIS PAGE LEFT INTENTIONALLY BLANK

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

ELAPSED TIME BETWEEN SEEKS

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TIME MSECS		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	REPORT 2
0	0	0.	0.	0-	0	I	
5935	5935	2.066	2.066	1-	1	Ixx	
7879	13814	4.809	2.743	2-	2	Ixxx	
10424	24238	8.436	3.629	3-	3	Ixxxx	
13957	38195	13.297	4.859	4-	4	Ixxxxx	
13675	51870	18.058	4.761	5-	5	Ixxxxx	
14260	66130	23.023	4.965	6-	6	Ixxxxx	
14658	80788	28.126	5.103	7-	7	Ixxxxx	
13735	94523	32.908	4.782	8-	8	Ixxxxx	
13192	107715	37.500	4.593	9-	9	Ixxxxx	
12161	119876	41.734	4.234	10-	10	Ixxxx	
11550	131426	45.755	4.021	11-	11	Ixxxx	
10956	142382	49.569	3.814	12-	12	Ixxxx	
10396	152778	53.189	3.619	13-	13	Ixxxx	
9580	162358	56.524	3.335	14-	14	Ixxx	
9380	171738	59.789	3.266	15-	15	Ixxx	
8894	180632	62.886	3.094	16-	16	Ixxx	
8319	188951	65.782	2.896	17-	17	Ixxx	
7945	196896	68.548	2.766	18-	18	Ixxx	
7182	204078	71.048	2.500	19-	19	Ixxx	
6454	210532	73.295	2.247	20-	20	Ixx	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
567	278696	97.026	0.197	49-	49	I	

287238 ENTRIES TOTAL AVERAGE = 16.89080 VARIANCE = 35200.476 STANDARD DEVIATION = 187.618

Figure 7-21. Elapsed Time Between Seeks Report

7.6 Default Option Alteration

Most users rely upon the standard MSM Report formats and their default values as these suit a wide range of needs. A capability to change the reports is built into MSMDRP. The general form for all option requests are as follows: The first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirement being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" entered in it. This card must be present whether or not any other input options are selected.

There is no specific order required of the options, and multiple entries of each are permissible. If several inputs refer to the same report, the last one encountered will have precedence. If a report is turned off by default and is modified, it will be turned on through the request for modification. The chart below shows the available actions: the mnemonic code for the user to identify the action; the function; and the default.

Mnemonic	Function	Default (indicated in parentheses)
AREA	Request file code references made to a specific area of a specific device	(not provided)
DEBUG	Debug	(no debug)
ERROR	Do not stop on Input Error	(stop)
FILDEF	Define system files by name	(no names used)
END	This card must be present	
MODULE	Produce the SSA Module Usage Report by Job	(no report produced)
NCONN	Process a limited number of connects	(total tape processed)
NREC	Process a limited number of tape records	(total tape processed)
OFF	Turn reports off	(all reports ON except reports 12,16, 18,20 - see table 7-1)
ON	Turn reports on	(all reports on except reports 12,16, 18,20 - see table 7-1)
PROJ	Produce the Connect Summary Report by Userid/SNUMB	(no report produced)

RN	This option must be selected when .MSMDRP is used to process WW6.4 data or when MSMDRP is executed on a WW6.4 system	(process WW7.2/4JS data on a WW7.2/4JS system)
TIME	Set a timespan for measurement	(no time criterion)
TIMEQ	Change time quantum for Chronological Device Utilization Report	(report is off - default value is 60 seconds)
USERID	Suppress userids from reports	(userids printed)
RATECH	Change time quantum for Connects Per 10 Minute Report	(report is off - default value is 10 minutes)
CAT	Turn on the Cat/File String Report	(report off)
RATE	Request the Connect Per 10 Minute Report for specific user jobs.	
LIMITS	Limit the amount of processing performed and reports produced.	
ZERO	Process jobs with an activity number of zero.	

7.6.1 Monitor a Specific Device Area (Action Code AREA). This option allows a user to specify specific areas of a device for which all jobs referencing this area are to be highlighted. The format of the display is that of a File Code Summary and contains those jobs and file codes that reference the area of interest.

The device to be investigated is identified via the PUB and IOM number. The specific areas of interest are identified as beginning at the starting address defined in llinks. The length of the area is also in llinks, with a zero meaning the end of the device. A total of ten possible areas are allowed. The format for this card is shown in figure 7-25.

See subsections 7.5.12 and 7.5.16 for complete details on the report format generated with this user option. This report is off by default and will be activated by the processing of this action code.

The following instructions are the procedures that should be followed when using this option:

- o Determine the device and area on the device to be examined. For purposes of this example, assume we want to analyze cylinders 300-420 on device 5, IOM-0, PUB-8, and cylinders 200 to the end of device 7, IOM-0, PUB-8. In addition, we want to report all connects issued to the specified area of device 5, but report accesses to specified area of device 7 only if they exceed 300 connects.

- o Convert the cylinder addresses to LLINKS. For purposes of this example, assume that we are analyzing 450 disk drives:

  cylinder 300 = (300x760) = 228000 sectors = (228000)/5 = 45600 LLINKS

  cylinder 200 = (200x760) = 152000 sectors = (152000)/5 = 30400 LLINKS

- o Determine the length of the search. For device 5, we want 120 cylinders, and for device 7, we want to go to the end of the device.

- o Convert the cylinder lengths to LLINKS.

  120 = (120x760) = 91200 sectors = 91200/5 = 18240 LLINKS

  The length for the second search is zero since we want to go to the end of the pack.

- o Fill out the cards following the format on figure 7-25:

```
AREA
  2
0 8 5 45600 18240  0
0 8 7 30400 0      300
```

7.6.2 System Debug (Action Code DEBUG). This is a restricted option for GMF system developers. DEBUG should only be used with guidance received by CCTC/C751. The option consists of the word DEBUG on the first card and any one of the following values on the second card:

9999 - perform trace logic checking

9998 - perform GEFSYE debug

9997 - perform accounting debug

9996 - perform SYSOUT debug

1-63 - debug this program number

-1 - de[illegible]g connects without a valid seek address

-N - [illegible] linkings of the FC array

7.6.3 Co[illegible] Reduction After an Input Option Error (Action Code ERROR). Th[illegible] allows data reduction to continue when an error has been detected and reported in an input option request. The default value reports the error and aborts the data reduction procedures. The format for this option is the word ERROR on the data card.

7.6.4 Specify System File Names (Action Code FILDEF). This option allows the user to specify the name of each system file displayed in the

Card 1 = A
Card 2 = N
Card 3 = B C D E F G

Where

A = The word AREA
N = The number of areas to be specified. A maximum of ten areas are permitted. A Card #3 must be present for each area requestred.
B = IOM number
C = Pub number
D = Device number
E = Starting address in llinks
F = Length of area in llinks
G = Do not report this file code unless number of connects exceeds this value

The following definitions apply to this option.

Device Type	Numbers Cylinders	Number Sectors/ Cylinder	Number Sectors/ Block (LLINK)
180	200	360	5
181	200	360	5
190	407	589	5
191	407	760	5
450	811	760	5
500	811	760	5
501	840	1280	5

Figure 7-25. Specific Device Area Report Card Input

System File Use Summary Report discussed in subsection 7.5.9. This option is specified with a set of data cards. The first data card contains the word FILDEF. The second data card contains the number of system files to be described on the following cards. The following cards each contain a single pair of data points separated by at least one blank. The first data point is the system file number and the second data point is the desired system file name.

The standard output of the System File Use Summary Report is to label each system file as System File 1, System File 2, etc., corresponding to GCOS-HI-USE, GCOS-LO-USE, etc. In order to know the correct order of the file names, the user should check the $ FILES section of the startup deck. The order of the files in the $ FILES section of the startup deck is the order they are referenced in the report.

7.6.5 <u>End Card (Action Code END)</u>. This card must be present at all times and must be the last data card supplied. It consists of the word END entered on the card.

7.6.6 <u>Produce the SSA Module Usage Report by Job (Action Code MODULE)</u>. This option allows the user to produce the SSA Module Usage Report. This report will list every SSA module used by every job that was run during the monitoring session. See subsection 7.5.11 for details concerning this report. This report is off by default and cannot be turned on by using the ON option. *This report can be activated only by entering MODULE on the data card.*

7.6.7 <u>Record Limitation by Connects (Action Code NCONN)</u>. This option allows a user to process only a specific number of connects. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NCONN with the second card containing the number of connects to be processed.

7.6.8 <u>Record Limitation by Records (Action Code NREC)</u>. This option allows a user to process only a specific number of tape records. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NREC with the second card containing the number of tape records to be processed.

7.6.9 <u>Turn a Report Off (Action Code OFF)</u>. This option allows a user to turn a report off that is on by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned off with this option. Two data cards are required to use this option. The first card contains the word OFF and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.10 Turn a Report On (Action Code ON). This option allows a user to turn a report on that is off by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned on with this option (#9,10,12,15,16,17,18,20). Two data cards are required to use this option. The first card contains the word ON and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.11 Produce Connect Summary Report by Userid/SNUMB (Action Code PROJ). This option allows the user to specify up to a total of 40 USERIDs and SNUMBs for which he wants the Connect Summary Report by Userid/SNUMB produced. The number of SNUMBs requested cannot exceed 10. In addition, the user can request the entire File Code Summary Report, or the user may want to see the File Code Summary Report only for a prespecified set of jobs or USERIDs, or the user may not want the File Code Summary Report at all. For example, the user can request 35 different USERIDs and 5 SNUMBs or 40 different USERIDs and 0 SNUMBs or 30 different USERIDs and 10 SNUMBs or 3 different USERIDs and 6 SNUMBs, etc. The format for this option is shown in figure 7-26. If values of zero are desired, they must be punched on the card. A blank is not equivalent to a zero. The Connect Summary Report will indicate for each requested USERID or SNUMB the number of connects made by the job or USERID. If a requested SNUMB also has a requested USERID, the number of connects issued by that job will be reported twice in the summary report. Refer to subsection 7.5.14 for a description of the report to be produced with this option. If the user desires to see a File Code Summary Report, it will be turned on via this option. The user does not need to use the "ON" input option.

7.6.12 Reduce WW6.4 Data or Process MSMDRP on a WW6.4 System (Action Code RN). This option requires two cards. The first card has the letters RN and the second card one of the following numbers:

1 - WW6.4/2H system processing WW6.4/2H data
2 - WW6.4/2H system processing WW7.2/4JS data
3 - WW7.2/4JS system processing WW6.4/2H data

The default is a WW7.2/4JS system with WW7.2/4JS data.

7.6.13 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours of which only a few may be of interest. This option allows a user to specify the timespan (or spans) to displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only from 0900 to 1700 in all reports.

If the entire reduction will have a set timespan, the name "TOTAL" is used. Histogram reports cannot be individually timespanned. All timespans of "other" reports will be bounded by the overall report timespan, if one will be used. Up to five timespans for each report type may be specified,

Card 1 = A
Card 2 = B C D
Card 3+ = E
Card 4+ = F

Where

A = The word *PROJ*

B = 0 if Connect Summary Report is desired, but no File Code Summary Report is desired. The 0 must be punched on the card. A blank is not equivalent to a 0.

B = 1 if Connect Summary Report is desired and a complete File Code Summary Report is wanted

B = 2 if Connect Summary Report is desired and only a partial File Code Summary Report is wanted

C = Number of Userids (30 MAX)

D = Number of SNUMBS (10 MAX)

E = A total of C Userids with one Userid per card. If C = 0, this card is not present.

F = A total of D SNUMBS separated by at least one blank space. All SNUMBS should fit on one card. If D = 0, this card is not present.

* - The values B, C and D must be separated by at least one blank column.

Figure 7-26. Limited File Code Summary Input Card Format

7.6.16 Change the Time Quantum Value for the Connect Per 10 Minute Report (Action Code RATECH). The user can change the time quantum value used to produce the Connect Per 10 Minute Report by inputting the quantum in seconds. Two cards are required. The first card contains the word RATECH and the second card contains the new quantum in minutes. The default value is 10 minutes.

7.6.17 Turn on the Cat/File String Report (Action Code CAT). This option, consisting of the word CAT on the data card, will turn on the Cat/File String Report (see subsection 7.5.13).

7.6.18 Request the Connect Per 10 Minute Report for Specific User Job (Action Code RATE). This option will allow the user to obtain the Connect Report for specific jobs as well as for the TimeSharing Subsystem and the Total System (see subsection 7.5.24). Card number 1 contains the word RATE, card number 2 the number of jobs desired (a maximum of 8 are permitted), and card number 3 the SNUMBs of the jobs desired. In addition to the requested jobs, the TimeSharing Subsystem as well as the Total System will also be reported. If multiple copies of TSS are in use, all activity will be reported under the single job name of TS1. If the user wants to obtain this report for only TimeSharing and the Total System, then he simply needs to use the "ON" input option using the name "RATE" for the required report ID.

7.6.19 Limit the Processing and Output (Action Code LIMITS). This option will allow the user to control the amount of output produced and the amount of record processing performed. Card number 1 contains the word LIMITS and card number 2 contains either the word ONLYSP, NOHIST, SUMARY, UTIL, or DEVICE. If the word ONLYSP is used then the Mass Store Monitor program will process only those data records that are generated by the SNUMBs requested under the RATE input option (see subsection 7.6.18). All other data will be ignored. The user must take care when examining the histograms and reports that are produced. The user must remember that only a limited amount of data has been processed. If the word NOHIST is used then no seek or space utilization histograms will be produced. This option can be used in conjunction with the ONLYSP option (must have two LIMITS input cards) or can be used by itself. In the latter case, all data will be analyzed, but no histograms will be produced. Finally, the user can request a summary of the seek movement activity. He can obtain this summary whether or not he selects to produce the set of histograms. To obtain the summary report, he must type SUMARY on a card immediately following the LIMITS card. A summary listing will not be produced for the space histograms, as this summary information is meaningless for this set of histograms. If the word UTIL is used, then only the Device Utilization by Device Number Report will be produced. This option will significantly reduce the run time of the MSMDRP and allow the user to determine those devices with high utilization. Once these devices have been identified, the user can run a second job using the DEVICE option. If the word DEVICE is used, the Data Reduction Program will analyze only the specific unique device IDs requested on the following card. If this option is used, the third card contains the number of unique devices to be checked (not to exceed 20) and the fourth card contains a list of unique device ID numbers obtained from the Device Utilization Histogram.

7.6.20 Zero Activity Processing (Action Code ZERO). This option allows the user to define up to 10 jobs which the user desires to see handled as normal activities, even though they process with an activity number of zero. Under normal conditions, any activity processing with an activity number of zero will be considered a System Scheduler Job (see subsection 7.5.15). To use this option, the first data card should have the word ZERO. The second data card contains the number of jobs following on the third card. This number may not exceed 10. The third card contains the list of SNUMBs, separated by at least one blank column.

7.7 JCL

The data reduction procedures consist of a single FORTRAN program having a main level and multiple subroutines.

A description of the more important JCL cards is presented below (see figure 7-28).

The $:LIMITS card should be studied to meet user needs. The run time (99) and output limit (30K) may both need to be altered as required by the duration of the monitoring run. The MSMDRP requires 73K of memory in order to execute plus an additional 2K for SSA space. During the initial loading process, MSMDRP will actually require 81K of memory, but 10K will be released immediately upon loading.

The statement:

```
$    DATA        I*
```

is used to identify the data cards that follow as described in subsection 7.6. At least one data card is required, that being an "END" request.

7.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

```
Col 1    8          16
$        IDENT      1820251/30/3044
$        SELECT     B29IDPXO/OBJECT/MSM
$        TAPE       01,X1D,,12345
$        LIMITS     99,73K,-4K,30K
$        DATA       I*

(The following is a suggested set of data cards)

ON
RATE
CAT
ON
CHRONO
LIMITS
SUMARY
END
$        ENDJOB
```

Figure 7-28. MSMDRP JCL

If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If the types in "N", then the program will be terminated and all reports will be produced.

7.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

```
$ XBAR     IOM-1,PUB-14,IOM-0,PUB-14
           IOM-2,PUB-14,IOM-2,PUB-13,
           IOM-0,PUB-13,IOM-1,PUB-13
```

<u>MPC Cards</u>

```
$ MPC-0    PSI-2,IOM-0,PUB-10,PSI-1,IOM-1,PUB-10,
           PSI-3,IOM-2,PUB-10,PSI-3,IOM-2,PUB-11,
           PSI-1,IOM-1,PUB-11,PSI-2,IOM-0,PUB-11

$ MPC-2    PSI-0,IOM-1,PUB-14,PSI-2,IOM-0,PUB-14
           PSI-1,IOM-2,PUB-14,PSI-1,IOM-2,PUB-13
           PSI-0,IOM-0,PUB-13,PSI-2,IOM-1,PUB-13
```

<u>Chart</u>

IOM/Channel	MPC/PSIA	IOM/Channel	MPC/PSIA
0-10	0-2	1-14	2-0
1-10	0-1	0-14	2-2
2-10	0-3 *	2-14	2-1
2-11	0-3	2-13	1-1 **
1-11	0-1	0-13	1-0
0-11	0-2	1-13	1-2

* problem 1
** problem 2

The problems described by the above procedures could be solved by redesigning the crossbar cards in the following manner:

```
$ XBAR     IOM-0,PUB-10,IOM-1,PUB-10,
           IOM-2,PUB-10,IOM-0,PUB-11,
           IOM-1,PUB-11,IOM-2,PUB-11

$ XBAR     IOM-1,PUB-14,IOM-0,PUB-14,
           IOM-2,PUB-14,IOM-1,PUB-13,
           IOM-0,PUB-13,IOM-2,PUB-13
```

If the I/O request cannot be granted, because either the channel or device being requested is currently busy, the request will be queued. This request will only be serviced when both a channel and device are free. When queuing occurs for a channel, GCOS will indicate the request queued over the primary channel. A primary channel is that channel which appears first on the $XBAR card, for a given string of devices. Therefore, all channel queue histograms are presented only for primary channels. However, a queue on the primary channel actually means that all channels, both physical and logical, connected to the desired device were busy. When the request is finally granted, a trace type 7 is issued.

A table is used to hold the device number, channel number, IOM number, I/O queue entry address, and the time the T22 trace event occurred. With the occurrence of each T22 event, the table entry is filled to mark the linking of the I/O requests. At this time, the required computations for determining the channel and device queue length are made. Channel queue histograms are produced for both tape and mass store devices, while device queue histograms are produced only for mass storage devices. The channel and device queue time also begins at this point and will be updated with the occurrence of the trace 7 event for this I/O request.

With the eventual occurrence of the trace 7 event for the I/O request, several updates are required to the common tables. The I/O queue time data are generated for the channel and device and collected for the appropriate histograms. It should be noted that it is possible for a device or channel to show no queuing, but yet they will display I/O queue time in their queue time histograms. The reason for this is that the queue time histogram is reporting the time difference between a T7 trace and a T22 trace. The T7 trace will not be issued until the actual I/O is initiated (i.e., a logical channel and the device become available). The following scenarios will attempt to illustrate this point:

Scenario 1

- Seven devices are configured over two logical channels.
- Device 1 receives a connect. At this point in time, device 1 is busy and 1 logical channel is busy.
- Device 2 receives a connect. At this point in time, devices [illegible] and 2 are both busy and both logical channels are busy.
- Device 3 receives an I/O request. Device 3 is not currently busy and therefore this I/O should be able to be initiated. However, there are no available channels and therefore the I/O must be delayed. Note that the queue length histogram for device 3 will report a 0 queue length, but may indicate a substantial queue time.
- Device 3 receives 3 additional I/O requests (4 outstanding requests exist at this point). The Channel Monitor will still report a queue length of 0 for device 3. The queue at device 3 exists not because of the device, but rather because of a channel shortage. The queue time histogram for device 3, however, will show a substantial queue time.
- Device 1 completes its I/O and device 3 initiates one of its waiting I/Os.
- Device 3 receives another I/O request.

- o At this point in time, device 3 is busy and has 3 additional requests queued for it. Therefore, the CMDRP will now report a queue length of three. Thus, we see that device queues may increase in a nonsequential fashion.

Scenario 2

- o A channel queue of length 0 exists when there is at least 1 nonbusy, primary/logical channel.

- o When all primary/logical channels become busy, the length of the channel queue is calculated by summing the length of all device queues on that channel.

- o As an example, assume we had two channels configured with three devices. If device 1 became busy and 4 more requests were made for that device, we would have a device queue of 4 and a channel queue of 0. If during the same time, device 2 received 3 requests, we would generate a queue length 2 for that device but the channel queue would still remain at 0. It should be noted, however, that even though no channel queueing is being reported, the channel queue time histograms will show significant queue times. This is because of the delay between the request for the I/0 (trace 22) and the actual connect (trace 7). In the situation described thus far, device contention, not a shortage of channels, is the problem. If we now assume that a connect was issued to device 3, we have now created a channel problem. There is no available channel for the I/0 request. At this point, a channel queue of 7 (4 requests for device 1 + 2 requests for device 2 + 1 request for device 3) would be generated. Thus, we see that channel queues may increase in a nonsequential fashion.

Even if a channel is available, upon the occurrence of a T22 trace, several milliseconds might pass before the system generates the T7 event. This is especially true on a very busy system. A connect queue entry is now filled with data to be used for the I/0 service time histograms. This connect queue holds the IOM, channel, device number, and the time of the trace 7 event. The channel and devices status table entries are also marked busy at this point. As a confidence test, the channel status is sensed at the start of processing for each trace 7 event for a nonbusy status. If it is busy, a lost interrupt is considered to have occurred since it is impossible for a connect to be issued to a busy device or channel. Device access histogram data and an IOM command execution count are also generated at trace 7 event time.

The next logical event for the I/0 process is the termination interrupt originated by the IOM at the I/0 data transfer completion. The signal for this event is transmitted by the IOM to the processor through the SCU as a request for the processor to service the I/O completion. The type 4 trace event contains the IOM number and channel for I/O termination. These data are used to determine the I/O service time by finding the time difference of the connect event and the terminate event. The time difference is collected and displayed in histogram form for each mass store and tape channel as well as for all mass store devices. The channel and device queue length are also adjusted at this point to reflect the absence of a queue being serviced for this channel and device.

It must be noted that exceptions to the normal I/O process are to be expected and must be accounted for in the reduction program. All the exceptions encountered so far have been diagnosed and coding in the program will allow for exceptions. Some of the exceptions include the following:

- o System programs that avoid the T22 trace by generating their own queue entry and by starting the connect immediately.
- o System programs that manipulate the I/O priority by linking themselves ahead of I/O requests already in the queue.
- o I/O requests for device zero (MPC) which makes a channel busy but not a device.
- o Lost interrupts from an I/O connect which leave the connect table and status table in an active state forever, if not detected.
- o Special controller commands which do not involve the channel or device.

8.4 Data Reduction Methodology

The CMDRP currently uses random I/O (File 58) to process histogram data. This feature allows the CMDRP to process an unlimited number of channels/devices with a minor increase in memory requirements. As delivered, the CMDRP will process 75 mass storage devices and 40 mass storage/tape channels. It will produce 106 unique histograms with a minimum amount of random I/O. If the number of channels or devices is insufficient, the user will need to edit file B29IDPXO/SOURCE/CM. The user should enter the edit subsystem and process the following command:

```
B RS:/NRDEVXX=XX75/;*:/NRDEVXX=XX New number of devices/

B RS:/NRCHANXX=XX40/;*:/NRCHANXX=XX New number of channels/
```

For each additional channel the size of the program will increase by 55 words and for each additional device the program will increase by 25 words. In the above edit the character "X" signifies a single space.

The last variable that will need to be changed is RPTCNT. This number represents the total number of histograms that will be processed with negligible random I/O. To calculate the total number of histograms that will be produced under your configuration, the following formula should be used.

(number of mass store devices)*3 + (number of tape and mass store channels - both logical and physical) + (number of tape and mass store physical channels)*2.

If this value is less than 106 no change is required. If the value computed is greater than 106, the user may alter this value. This will help to decrease CPU/IO time but will increase storage by 80 words for each increment above 106. This tradeoff between CPU/IO time and memory must be made at the discretion of the user. In order to change this value, the following edit function should be performed:

```
B RS:/RPTCNTX=106/;*:/RPTCNTX=New value
```

As in the earlier edit example, the character "X" should not be typed, but is being used to represent a blank column. After performing the above edits, the user should recompile the source program by entering the card subsystem and issuing a run command.

8.5 CMDRP Output

Reports generated from the Channel Monitor may vary from one collection period to another due to the difference in configuration of the hardware. Report numbers are preassigned to the histogram reports which are hardware independent and are dynamically assigned to histograms which denote the channel and device uniquely configured on each IOM.

The following subsections will describe all reports produced by the CMDRP and subsection 8.6 will describe the user input options.

8.5.1 System Configuration and Channel Usage Report (File 57). This report documents the system identification, configuration, and the date and time of the monitoring period, as well as reporting the usage of all configured I/O channels. Figure 8-2 is an example of this report. The heading line indicates the software version number that corresponds to this document. The version number should be 09-83 CHG-7. The first line after the heading provides the tape number(s) the report was generated from, the system identification, the date (in the form year, month, and day - YYMMDD), and the start and stop times (HH:MM:SS) of the MONITORING SESSION. The next several lines of output describe the overhead of all GMF monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM".

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

```
******************************************************
* * *          C H A N N E L   M O N I T O R     * * *
* * *             VERSION 09-83  CHG-7           * * *
******************************************************

          TAPE # 00860
      SYSID              DATE      START TIME   STOP TIME
     AZ-DP3,DPS1.3     83-08-04     10:47:24    11:16:28      FOR A TOTAL OF 0.48 HOURS

  MONITOR OVERHEAD - MONITOR         TIME(SEC)        % OVERHEAD
                     EXEC               79              1.51
                     MSM                79              1.51
                     CM                 57              1.09
                     NAME                0              0.
                     FMS                 0              0.
                     TOTAL                              4.11

CONFIGURATION:  TRIPLE PROCESSOR DPS4/5,  DUAL  IOM,  1536K  MEMORY - 52 OF WHICH WERE HCM
                  THE HCM DOES NOT INCLUDE 3K FOR .CALC AND 6-8K FOR FILSYS
  PROCESSOR OPTIONS - 8K CACHE, EIS OPTIONS   INSTALLED, FAST MEMORY ACCESS INSTRUCTION OVERLAP ON
                 CPUS ACTUALLY CONFIGURED = 3.00  CPUS ACTUALLY AVAILABLE = 2.50
                                   IOM NUMBER 0

       CHANNEL           TYPE          CROSSBAR       CONNECTS
        0-08            .DS450           0-09          13137
                                         0-10
                                         0-11
                                         1-08
                                         1-09
                                         1-10
                                         1-11
```

Figure 8-2. System Configuration and Channel Usage Report (Part 1 of 2)

0-09	.DS450	SEE ABOVE	13131
0-10	.DS450	SEE ABOVE	13129
0-11	.DS450	SEE ABOVE	13135
0-12	.DS191	0-13	189
		1-12	
		1-13	
0-13	.DS191	SEE ABOVE	188
0-14	.DMTM9	0-15	12148
		1-14	
		1-15	
0-15	.DMTM9	SEE ABOVE	12737
0-16	.DS500	1-17	40043
0-17	.DS500	1-16	9405

IOM NUMBER 1

CHANNEL	TYPE	CROSSBAR	CONNECTS
1-08	.DS450	SEE ABOVE	13140
1-09	.DS450	SEE ABOVE	13130
1-10	.DS450	SEE ABOVE	13136
1-11	.DS450	SEE ABOVE	13134
1-12	.DS191	SEE ABOVE	188
1-13	.DS191	SEE ABOVE	190
1-14	.DMTM9	SEE ABOVE	12147
1-15	.DMTM9	SEE ABOVE	4715
1-16	.DS500	SEE ABOVE	9405
1-17	.DS500	SEE ABOVE	40353
1-22	.DMTM9	1-23	7457
1-23	.DMTM9	SEE ABOVE	7459

Figure 8-2. (Part 2 of 2)

THIS PAGE LEFT INTENTIONALLY BLANK

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as:

$$\frac{\text{(CPU TIME Used by Monitor)}}{\text{(TOTAL Elapsed Time)x(Number of Processors)}}$$

Following the overhead description are three lines of configuration information describing the number of processors, IOMs, and amount of memory configured to the system. In addition, the size of GCOS Hard Core, the size of the Core Allocator and the size of FILSYS is also presented. The third line of the configuration data indicates the number of processors actually configured and actually available. These numbers might be different than shown on the first line due to the assigning and releasing of processors. In figure 8-2, we see that one processor was released for a period of time (i.e., CPUs actually available is equal to 2.50). The actual time that processors were available or released is indicated in the status message printouts (see subsection 8.5.15).

The next portion of the report documents the channel configuration by IOM, listing each configured channel number (both tape and disk), the device type configured to that channel, and the channel crossbarring. The crossbar column shows those channels that are crossbarred to the channel identified under the channel column. If SEE ABOVE is found, the crossbarring has been displayed on a preceding channel. The I-CC format of each channel description identifies the IOM and the channel number being discussed. The last column of this report displays the number of all connect types issued over that channel. This section will be repeated for each IOM configured to the system. This report is always generated and cannot be turned off.

8.5.2 System Summary Report (File 57). The System Configuration and Channel Usage Report and the System Summary Report may be used to assess overall system utilization. Figure 8-3 is an example of the System Summary Report. The first set of lines shows the number of connects to the monitored mass storage subsystems compared to the total connects issued (TAPE+DISK) and the connect rate per hour over the subsystem. Most systems will show a small number of Control Connects being generated by the MPCs configured to the system. These Control Connects will be summed together and listed as a separate subsystem line. Analysis on a Shared Mass Storage System shows the number of MPC connects generated to be a significant percentage of the total connects generated. The next lines show the breakdown of the mass storage and tape connects by the IOM over which they were issued. The final part of this report is a list of the commands (octal code and mnemonic) issued to the mass storage subsystem and the count of each issued during the monitoring session. This report is always generated and cannot be turned off.

```
*************************************************************************
  * * *          S Y S T E M   S U M M A R Y   R E P O R T          * * *
*************************************************************************

TOTAL CONNECTS TO DSS191         755 OF    261697 ( 0%) AT A RATE OF      1558. PER HOUR
TOTAL CONNECTS TO MS0450      105072 OF    261697 (40%) AT A RATE OF    216805. PER HOUR
TOTAL CONNECTS TO MS0500       99206 OF    261697 (37%) AT A RATE OF    204701. PER HOUR
TOTAL CONNECTS TO IOM-0       127242 OF    261697 (49%) AT A RATE OF    265087. PER HOUR
TOTAL CONNECTS TO IOM-1       134455 OF    261697 (51%) AT A RATE OF    280114. PER HOUR

                                 COMMAND                 COUNT
                              23 RD ASC                      8
                              25 READ                   160111
                              31 WRITE                   41778
                              32 WR ASC                      3
                              33 WR-VER                   3133
                              TOTAL                     205033
```

Figure 8-3. System Summary Report

A well performing system, under a heavy workload, should show a high utilization of the configured resources. Figure 8-3 shows that the I/O activity is evenly divided between the MS0450 and MS0500 subsystems. All signs indicate that if system throughput is being affected by disk activity, then the MS0450s and MS0500s would be the probable cause of such problems.

The next item to check on these two reports should be the channel usage. The two highest used logical channels of any subsystem should be on a separate PSIA channel of a two-PSIA channel subsystem (see subsection 8.3). Referring to figure 8-2, one can see that logical channel 8 of IOM 0 and IOM 1 has the highest usage, and this is the proper configuration (refer to detailed description in subsection 8.3). If the highest used logical channels are not on separate PSIA channels, the $ XBAR card in the startup configuration section is suspected as the cause. The channels are used in the order given on the $ XBAR card (i.e., if the primary channel is busy, the next channel tried is given on the crossbar). The alternate use of PSIA channels for maximum simultaneity must, therefore, be appropriately specified in the boot deck.

While looking at the System Summary Report, it is also of interest to note the ratio of READ commands to WRITE commands (nearly four to one in this example). This gives an indication of the nature of the usage of the mass storage space. A quick look at the number of write/verify (WR-VER) commands executed is also of interest as they are essentially double (WRITE, then READ) data transfer commands which require more device and channel time.

The general fraction of utilization for each logical channel gives an indication of the degree of simultaneity of access to the subsystem. If only N of the configured logical channels have nonzero counts, then there were never more than N accesses being performed simultaneously by the subsystem. The proportional relationships among the counts of accesses made over each of the logical channels are quantitative indications of the frequency of occurrence of specific levels of simultaneity. As an example, if we look at figure 8-2, we see that all of the logical channels for a given disk subsystem are receiving nearly an identical number of connects. This would indicate that a significant number of times all of the logical channels were busy and the degree of simultaneity was very high. In this example, channel queuing (i.e., shortage of channel power) could very well be a problem. This is not to infer that device queuing is not also a problem, just that channel queuing may be a problem. If the number of accesses to the lowest priority channel is a smaller percentage of the total accesses, then channel queuing will probably not need to be examined.

8.5.3 System Traces Captured by Monitor Report (File 57). This report contains the number of occurrences of each specific trace type recorded on the data collector tape processed by the CMDRP (figure 8-4). This report provides little, if any, information required by the user for his analysis. This report is always generated and cannot be turned off.

8.5.4 Channel Status Changes Report (File 57). This report lists the initial status for all tape and disk channels configured to the system (figure 8-5). If, during the course of the monitoring session, a given channel or IOM was dropped or added to the system (dynamic reconfiguration) a new report will be produced indicating the activation or deactivation changes and the time that the change occurred. Finally, this report will indicate whether the SSA cache option and FMS cache option are active, and if so, will indicate their initial status and any changes that occur to that status. If a given option is not active, a zero will be reported for each of the values. This report is always generated and cannot be turned off.

8.5.5 Physical Device, Device ID Correlation Table (File 57). Each mass storage device configured in the system is listed with a unique device ID. A typical report is presented in figure 8-6. This unique device is needed since different devices can have the same device number on the Honeywell 6000. (See Device ID 24, Device ID 33, and Device ID 49 of figure 8-6). These unique numbers are referenced in several reports produced by the CMDRP. This report is always generated and cannot be turned off.

8.5.6 Channel Statistics Report (File 57). The Channel Statistics Report is actually a series of reports used to summarize the queuing that occurred over the channels and devices. These reports are processed as a group and are produced by default. The entire series of reports can be turned off via the use of the "OFF" input option (see subsection 8.6.5).

8.5.6.1 Channel Busy and Device Busy Report. The Channel Busy and Device Busy Report is given in figure 8-7. These data are collected into an array during execution of the reduction program and indicates the number of times channel K and device N were both busy (connects had been issued and were currently in processing) at the I/O request link time. Remember that this report is presented by primary channel. The primary channel is busy when all logical channels associated with the primary channel are busy. As an example, figure 8-7 would indicate a potential channel shortage for channel 16 IOM 0 and a potential device contention problem for devices 9, 11 and 13.

8.5.6.2 Channel Busy and Device Free Report. The Channel Busy and Device Free Report, figure 8-8, is generated in a similar method to the Channel Busy and Device Busy Report. If a channel is busy a sufficient number of times, this will indicate the need for more channel power. If, for example, there are no more IOM ports available, a large amount of channel queuing can be a strong indication of the need for another IOM. An entry is made to this report if a connect to a given device is delayed because there are no available logical channels and there is no active connect being processed for the given device. Figure 8-8 shows severe channel contention for channel 16 IOM 0.

8.5.6.3 Channel Free and Device Busy Report. The Channel Free and Device Busy Report, figure 8-9, is generated in a similar manner to the previous

THIS PAGE LEFT INTENTIONALLY BLANK

SYSTEM TRACES CAPTURED BY MONITOR REPORT

TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT
00	0	01	0	02	0	03	0	04	3214	05	0	06	0	07	3214
10	129	11	0	12	0	13	0	14	636	15	0	16	0	17	0
20	0	21	450	22	3214	23	0	24	0	25	0	26	0	27	0
30	0	31	0	32	0	33	0	34	0	35	0	36	0	37	0
40	0	41	0	42	0	43	0	44	0	45	0	46	15	47	0
50	21	51	0	52	0	53	0	54	0	55	0	56	0	57	0
60	0	61	0	62	0	63	0	64	0	65	0	66	0	67	0
70	0	71	0	72	0	73	0	74	0	75	0	76	0	77	0

Figure 8-4. System Traces Captured by Monitor Report

CHANNEL STATUS CHANGES REPORT FOR NMCC2 ON 80/09/20

IOM	CHANNEL	DEACTIVATE/ACTIVATE CHANGE	TIME
0	08	INITIAL ACTIVE	16:50:49.7
0	09	INITIAL ACTIVE	16:50:49.7
0	12	INITIAL ACTIVE	16:50:49.7
0	13	INITIAL ACTIVE	16:50:49.7
0	14	INITIAL ACTIVE	16:50:49.7
0	15	INITIAL ACTIVE	16:50:49.7
0	16	INITIAL ACTIVE	16:50:49.7
0	17	INITIAL ACTIVE	16:50:49.7
0	18	INITIAL ACTIVE	16:50:49.7
1	08	INITIAL ACTIVE	16:50:49.7
1	09	INITIAL ACTIVE	16:50:49.7
1	12	INITIAL ACTIVE	16:50:49.7
1	13	INITIAL ACTIVE	16:50:49.7
1	14	INITIAL ACTIVE	16:50:49.7
1	15	INITIAL ACTIVE	16:50:49.7
1	16	INITIAL ACTIVE	16:50:49.7
1	17	INITIAL ACTIVE	16:50:49.7

INITIAL VALUES FOR SSA CACHE - LAL, MBA, SIZE
000000200000 000000001566 10

INITIAL VALUES FOR FMS CACHE - ABS ADDR, MBA, OPTION WORD, #320 WORD BUFFERS
000000000000 000000000000 000000000000 0

Figure 8-5. Channel Status Changes Report

THE PHYSICAL DEVICE, DEVICE ID CORRELATION TABLE

```
DEVICE ID - 1 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #06
DEVICE ID - 2 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #07
DEVICE ID - 3 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #08
DEVICE ID - 4 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #09
DEVICE ID - 5 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #10
DEVICE ID - 6 IS FOUND ON IOM-O PUB-08 AND IS DEVICE #11
             .
             .
             .
DEVICE ID -24 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #01
DEVICE ID -25 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #02
DEVICE ID -26 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #03
DEVICE ID -27 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #04
DEVICE ID -28 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #05
DEVICE ID -29 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #21
DEVICE ID -30 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #22
DEVICE ID -31 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #23
DEVICE ID -32 IS FOUND ON IOM-O PUB-12 AND IS DEVICE #24
DEVICE ID -33 IS FOUND ON IOM-O PUB-16 AND IS DEVICE #01
DEVICE ID -34 IS FOUND ON IOM-O PUB-16 AND IS DEVICE #02
DEVICE ID -35 IS FOUND ON IOM-O PUB-16 AND IS DEVICE #03
             .
             .
             .
DEVICE ID -49 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #01
DEVICE ID -50 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #02
DEVICE ID -51 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #03
DEVICE ID -52 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #04
DEVICE ID -53 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #05
DEVICE ID -54 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #06
             .
             .
             .
DEVICE ID -63 IS FOUND ON IOM-O PUB-17 AND IS DEVICE #15
```

Figure 8-6. Physical Device, Device ID Correlation Table

CHANNEL BUSY AND DEVICE BUSY REPORT FOR AZ-DP3 ON 83-08-04

```
CHANNEL  08  IOM  0  AND DEVICE 09 BOTH BUSY      2  TIMES    0.03 PERCENT
CHANNEL  08  IOM  0  AND DEVICE 11 BOTH BUSY      7  TIMES    0.07 PERCENT
CHANNEL  08  IOM  0  AND DEVICE 12 BOTH BUSY      4  TIMES    0.05 PERCENT
CHANNEL  08  IOM  0  AND DEVICE 14 BOTH BUSY      7  TIMES    0.07 PERCENT
          .
          .
          .
CHANNEL  16  IOM  0  AND DEVICE 09 BOTH BUSY   1584  TIMES   11.61 PERCENT
CHANNEL  16  IOM  0  AND DEVICE 11 BOTH BUSY   1751  TIMES    9.36 PERCENT
CHANNEL  16  IOM  0  AND DEVICE 12 BOTH BUSY     96  TIMES    0.44 PERCENT
CHANNEL  16  IOM  0  AND DEVICE 13 BOTH BUSY   1561  TIMES   15.47 PERCENT
CHANNEL  17  IOM  0  AND DEVICE 07 BOTH BUSY     45  TIMES    1.96 PERCENT
          .
          .
          .
```

Figure 8-7. Channel Busy and Device Busy Report

CHANNEL BUSY AND DEVICE FREE REPORT FOR AZ-DP3 ON 83-08-04

CHANNEL	08	IOM	0	BUSY AND DEVICE	07	NOT BUSY	7	TIMES	0.08	PERCENT
CHANNEL	08	IOM	0	BUSY AND DEVICE	09	NOT BUSY	2	TIMES	0.03	PERCENT
CHANNEL	08	IOM	0	BUSY AND DEVICE	11	NOT BUSY	7	TIMES	0.07	PERCENT
	.									
	.									
	.									
CHANNEL	16	IOM	0	BUSY AND DEVICE	01	NOT BUSY	4626	TIMES	78.55	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	02	NOT BUSY	4130	TIMES	81.96	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	03	NOT BUSY	2311	TIMES	82.07	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	04	NOT BUSY	1069	TIMES	82.61	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	05	NOT BUSY	19	TIMES	79.17	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	06	NOT BUSY	20	TIMES	83.33	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	08	NOT BUSY	224	TIMES	85.50	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	09	NOT BUSY	9289	TIMES	68.08	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	10	NOT BUSY	24	TIMES	72.73	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	11	NOT BUSY	10570	TIMES	56.53	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	12	NOT BUSY	11206	TIMES	51.35	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	13	NOT BUSY	7546	TIMES	74.79	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	14	NOT BUSY	612	TIMES	86.20	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	15	NOT BUSY	30	TIMES	83.33	PERCENT
CHANNEL	16	IOM	0	BUSY AND DEVICE	16	NOT BUSY	11	TIMES	91.67	PERCENT
CHANNEL	17	IOM	0	BUSY AND DEVICE	02	NOT BUSY	9	TIMES	8.57	PERCENT
CHANNEL	17	IOM	0	BUSY AND DEVICE	03	NOT BUSY	1	TIMES	11.11	PERCENT
	.									
	.									
	.									

Figure 8-8. Channel Busy and Device Free Report

CHANNEL FREE AND DEVICE BUSY REPORT FOR AZ-DP3 ON 83-08-04

CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	09	BUSY	734	TIMES	11.15	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	11	BUSY	1143	TIMES	11.13	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	12	BUSY	1618	TIMES	19.37	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	14	BUSY	2130	TIMES	22.51	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	16	BUSY	4720	TIMES	34.72	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	17	BUSY	1178	TIMES	14.67	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	18	BUSY	1572	TIMES	15.87	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	19	BUSY	996	TIMES	15.77	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	20	BUSY	162	TIMES	5.72	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	27	BUSY	2834	TIMES	25.01	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	30	BUSY	574	TIMES	10.65	PERCENT
CHANNEL	08	IOM	0	NOT BUSY AND DEVICE	31	BUSY	8	TIMES	1.24	PERCENT
CHANNEL	16	IOM	0	NOT BUSY AND DEVICE	01	BUSY	5	TIMES	0.08	PERCENT
CHANNEL	16	IOM	0	NOT BUSY AND DEVICE	04	BUSY	2	TIMES	0.15	PERCENT
	•									
	•									
	•									
CHANNEL	17	IOM	0	NOT BUSY AND DEVICE	10	BUSY	1342	TIMES	12.53	PERCENT
CHANNEL	17	IOM	0	NOT BUSY AND DEVICE	11	BUSY	93	TIMES	2.59	PERCENT
CHANNEL	17	IOM	0	NOT BUSY AND DEVICE	13	BUSY	3	TIMES	0.62	PERCENT

Figure 8-9. Channel Free and Device Busy Report

two reports. It is an indication of the number of times that an I/O request was queued because the device was busy but there were available channels. Significant values in this report would be an indication that possible relocation of files is required. The Mass Store Monitor can be used with this data tape to determine the files being accessed on this device. Figure 8-9 shows several devices on channel 8 IOM 0 with significant contention.

8.5.6.4 Channel Free and Device Free Report. The Channel Free and Device Free Report, figure 8-10, is an indication of the number of I/O requests that were executed by IOS immediately, without any queuing. This report is used as a backup to the previous reports. As can be seen in figure 8-10, several devices on channel 8 IOM 0 (14, 16, 17, 18, 19) have limited contention in that between 15-35 percent of all connects are delayed. This should be confirmed by the other reports (see figure 8-9). In addition, this figure reconfirms the previous discovery that channel 16 IOM 0 has serious contention problems in that only a small percentage of connects are issued without any delay.

8.5.6.5 GEPR Connect Report. The GEPR Connect Report, figure 8-11, shows the number of times that a trace type 7 was processed, without a preceding trace type 22. This unexpected sequence of traces is supposed to occur whenever the system processes a GEPR (I/O error) event. During data collection, the CM captures an I/O status word indicator and the data reduction program checks this status word to verify that a GEPR has actually occurred. This report will indicate how many confirmed GEPRs (i.e., the status word was set) have occurred and how many suspected GEPRs have occurred (i.e., the status word was not set). This report will be of little aid to the analyst and CCTC is still investigating the reason for this trace occurrence, when the status word is not set.

8.5.6.6 Lost Interrupt Report. The Lost Interrupt Report, figure 8-12, indicates that a trace type 7 is being processed for a busy device and/or channel. This is an impossible event and is an indication that a trace type 4 has not been generated. This is usually an indication that the system has generated a lost interrupt. However, if lost data has been generated during data collection, this report may indicate many lost interrupts, which really did not occur.

8.5.6.7 Device ID STIOS Not Connected Report. The Device ID STIOS Not Connected Report, figure 8-13, shows the number of start I/Os that were dropped for each device. A start I/O is dropped when it is found that a STIO trace has occurred for a device and within a user-defined timeframe (5-second default), no connect has been received for the start. The cause of this condition currently is under analysis. This report is also presented by the device ID and must be correlated by using the Device ID Correlation Report.

8.5.6.8 Entries Still in Queue Report. The Entries Still in Queue Report, figure 8-14, shows all entries remaining in the CMDRP queues. This report shows the device number, channel number, IOM number, queue location, and time (in milliseconds) the entry has been in the queue. These entries were active when the monitor terminated.

8.5.6.9 Device Free But Has a Queue Report. The Device Free But Has a Queue Report, figure 8-14.1, shows the number of times a connect was made to a device, the device was free (no active connect being processed), but yet there were outstanding requests waiting for the device. In most circumstances, this event would be an example of a device being delayed because of the lack of sufficient channel power (see subsection 8.3). As can be seen in figure 8-14.1, device 11 on channel 16 IOM 0 had this occurrence for 19.9 percent (3728 times) of all its connects. Of the 3728 times this occurred, 3444 times there were no available channels.

However, in figure 8-14.1, several devices report this occurrence, but yet indicate that there were free channels. Under normal circumstances, if there is excess channel capacity and outstanding I/O for a device, then that device should begin to process the outstanding I/O. The occurrence of this condition is an indication of a backlog of I/O requests that the GCOS system is unable to keep pace with. The channel or device have probably just been freed and GCOS has not yet been able to process the next outstanding request for that device.

8.5.6.10 500 Disk Drive Report. The 500-type disk drives are actually 2 logical devices which are configured as a single physical drive. Therefore, it is possible for a connect to be made to a 500 disk, the disk to be free, but yet the connect cannot be issued because the partner 500 disk (the second logical device) is busy. The 500 drives are configured as odd/even pairs. Figure 8-14.2 describes the number of times this event occurs. As can be seen in this figure, device 11 on channel 16 issued 16762 connects (89.6 percent of all connects issued to that device) that were delayed because its partner device (device 12) was busy. The average accumulated queue size for both devices when this event occurred, was 1.3. In a similar manner, 14,715 connects (67.4 percent of all connects issued to device 12) were delayed because its partner device (device 11) was busy. This is an excellent example of device contention between two different disk packs that are physically configured on a single drive.

8.5.7 Idle Monitor Report (File 57). If the Idle Monitor was active when the Channel Monitor was running an Idle Report will be produced next (see figure 8-15). The first few lines will indicate, for each processor, the

THIS PAGE LEFT INTENTIONALLY BLANK

CHANNEL FREE AND DEVICE FREE REPORT FOR AZ-DP3 ON 83-08-04

CHANNEL	08	IOM	0	AND	DEVICE	06	FREE	24	TIMES	100.00	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	07	FREE	8729	TIMES	99.92	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	08	FREE	24	TIMES	100.00	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	09	FREE	5846	TIMES	88.79	PERCENT
	.										
	.										
	.										
CHANNEL	08	IOM	0	AND	DEVICE	14	FREE	7325	TIMES	77.41	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	16	FREE	8863	TIMES	65.20	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	17	FREE	6844	TIMES	85.25	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	18	FREE	8330	TIMES	84.12	PERCENT
CHANNEL	08	IOM	0	AND	DEVICE	19	FREE	5319	TIMES	84.20	PERCENT
	.										
	.										
	.										
CHANNEL	16	IOM	0	AND	DEVICE	01	FREE	1241	TIMES	21.07	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	02	FREE	916	TIMES	18.18	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	03	FREE	477	TIMES	16.94	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	04	FREE	139	TIMES	10.74	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	05	FREE	5	TIMES	20.83	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	06	FREE	4	TIMES	16.67	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	08	FREE	39	TIMES	14.89	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	09	FREE	2499	TIMES	18.31	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	10	FREE	9	TIMES	27.27	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	11	FREE	6320	TIMES	33.80	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	12	FREE	10530	TIMES	48.25	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	13	FREE	966	TIMES	9.57	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	14	FREE	99	TIMES	13.94	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	15	FREE	6	TIMES	16.67	PERCENT
CHANNEL	16	IOM	0	AND	DEVICE	16	FREE	1	TIMES	8.33	PERCENT
	.										
	.										
	.										

Figure 8-10. Channel Free and Device Free Report

```
GEPR CONNECT REPORT FOR AZ-DP3 ON 83-08-04

CHANNEL  08  IOM  0  DEVICE  07  # OF CONFIRMED GEPRS    0 # OF SUSPECTED GEPRS    13
CHANNEL  08  IOM  0  DEVICE  09  # OF CONFIRMED GEPRS    0 # OF SUSPECTED GEPRS     9
CHANNEL  08  IOM  0  DEVICE  11  # OF CONFIRMED GEPRS    0 # OF SUSPECTED GEPRS    19
CHANNEL  12  IOM  0  DEVICE  15  # OF CONFIRMED GEPRS    0 # OF SUSPECTED GEPRS     5
```

Figure 8-11. GEPR Connect Report

```
LOST INTERRUPT REPORT FOR AZ-DP3 ON 83-08-04

IOM  0  PUB  8  LOST INTERRUPT COUNT     1
```

Figure 8-12. Lost Interrupt Report

```
DEVICE ID STIOS NOT CONNECTED FOR AZ-DP3 ON 83-08-04

DEVICE ID   8  #  OF STARTS DROPPED     1

DEVICE ID  10  #  OF STARTS DROPPED     1
```

Figure 8-13. Device ID STIOS Not Connected

ENTRIES STILL IN QUEUE FOR AZ-DP3 ON 83-08-04

DEVICE #	CHA #	IOM	QUEUE	TIME IN QUEUE(MS)
44	16	0	185665173	0
7	8	0	563676821	4
21	8	0	644416171	10
13	8	0	666437269	21
7	8	0	644416160	39
11	8	0	234952330	44
6	8	0	372310016	46

Figure 8-14. Entries Still in Queue

DEVICE FREE BUT HAS A QUEUE REPORT FOR AZ-DP3 ON 83-08-04

PUB	IOM	DEVICE	TIMES FREE WITH QUEUE	PERCENT	AVG QUEUE SIZE	TIMES CHNL BUSY	TIMES CHNL FREE
08	0	09	5	0.08	1	0	5
08	0	10	2	0.80	1	0	2
08	0	11	81	0.79	1	0	81
08	0	12	366	4.38	1	3	363
08	0	14	27	0.29	1	0	27
08	0	16	153	1.13	1	0	153
08	0	17	8	0.10	1	0	8
08	0	18	12	0.12	1	0	12
08	0	19	50	0.79	1	0	50
08	0	20	190	6.71	2	0	190
08	0	27	57	0.50	1	0	57
08	0	30	498	9.24	1	0	498
08	0	31	9	1.40	1	0	9
12	0	23	82	13.42	1	0	82
16	0	01	136	2.31	1	131	5
16	0	03	19	0.67	1	18	1
16	0	04	79	6.11	1	77	2
16	0	08	2	0.76	1	2	0
16	0	09	813	5.96	1	718	95
16	0	10	10	30.30	1	5	5
16	0	11	3728	19.94	1	3444	284
16	0	12	65	0.30	1	48	17
16	0	13	1672	16.57	1	1598	74
17	0	02	37	35.24	1	7	30
17	0	08	195	42.58	1	2	193
17	0	10	4063	37.94	2	29	4034
17	0	13	59	12.22	1	5	54

Figure 8-14.1. Device Free But Has a Queue Report

500 DISK DRIVE REPORT FOR AZ-DP3 ON 83-08-04

PJB	IOM	DEVICE	TIMES THE PARTNER DEVICE WAS BUSY	PERCENT	AVG ACCUMULATED QUEUE SIZE
16	0	01	58	0.98	1.00
16	0	02	59	1.17	1.00
16	0	03	89	3.16	1.11
16	0	04	90	6.96	1.10
16	0	09	11	0.08	1.27
16	0	10	8	24.24	1.13
16	0	11	16762	89.64	1.31
16	0	12	14715	67.43	1.36
16	0	13	95	0.94	1.37
16	0	14	107	15.07	1.31
17	0	07	2	0.09	1.50
17	0	08	2	0.44	1.50
17	0	09	217	38.34	1.96
17	0	10	229	2.14	1.94
17	0	11	3	0.08	1.00
17	0	12	3	11.54	1.00
17	0	13	5	1.04	1.00
17	0	14	7	1.35	1.00

Figure 8-14.2. 500 Disk Drive Report

```
                    IDLE REPORT FOR SYSTEM OSCC2 ON 79-07-13

          PROCESSOR 0 WENT IDLE  101934 TIMES FOR A % IDLE OF  38

          PROCESSOR 1 WENT IDLE   64039 TIMES FOR A % IDLE OF  29
                        AVERAGE SYSTEM % IDLE WAS  34
                        % OF IDLE TIME DURING WHICH IO WAS ACTIVE  100
                        AVERAGE NUMBER OF OUTSTANDING IO's WHEN THE SYSTEM WENT IDLE WAS  7
```

DEVID	% OF IDLE TIME WITH ACTIVE IO	AVERAGE QUEUE SIZE
1	95	4.66
2	12	0.19
3	21	0.33
4	46	1.16
8	3	0.05
10	1	0.02
11	1	0.02
14	71	0.76
15	4	0.08
16	1	0.02
24	1	0.02

Figure 8-15. Idle Report

number of times that processor went idle and the percent idleness of the that processor. This is followed by an average idle percentage for the entire system. The next line of output indicates for what percentage of system idle time there was active disk I/O; i.e., I/O in progress or I/O request queued. This figure would give a good indication as to whether the CPU is going idle because of a lack of work or rather because work is being delayed by the slower peripheral devices. In figure 8-15, we see that for 100 percent of the idle time, the CPU was actually waiting on disk I/O to be performed. In this case, the CPU could be put to better utilization, if only the speed of the disk subsystem could be increased. The next line of output indicates how many outstanding I/Os were present when the system went idle.

Following this output, a table is generated for every disk device that had any active I/O on it when the CPU went idle. For each such device, the percent of Idle Time during which it had active I/O and the average queue size at that time is given. In figure 8-15, we see that during 95 percent of CPU idle time device ID #1 had outstanding I/O to an average length of 4.66. Device #4 had active I/O for 46 percent of the CPU idle time with an average length of 1.2. By examining this report, we can see that if the queues on these two devices could be reduced there is sufficient CPU power available to handle the additional workload that would be generated. The Device ID Correlation Report should be used to convert a unique device ID into an IOM, PUB, device number. This report was generated from a different monitoring session than all previous reports described thus far.

8.5.8 Proportionate Device Utilization Report (File 57). This report shows the proportionate utilization of each device configured on the mass storage subsystem. Figure 8-16 is an example. This histogram identifies each unique device ID (device number zero is an MPC controller) and provides both a count of the number of accesses made to each device (under the column headed INDIV. NUMBER) as well as the percent of all accesses which were to each device (under the column headed INDIV. PRC). The histogram shows the proportionate utilization of each device (i.e., the percent of all accesses which went to each device) in a graphical form. The physical device that each "Device ID" of the histogram represents is shown in the Physical Device ID Correlation Table (see figure 8-6). This report is always generated and cannot be turned off. In this report the user is looking for a device or devices which have significantly more utilization than others in the system. This highly used device would then be a potential bottleneck.

It is desirable, but not always practical, to have equal utilization for each device. The user should be reminded that data in figure 8-16 is cumulative over the monitoring period. The actual accessing pattern could have been periodic with the following form: Many accesses to device 4 followed by many accesses to device 3 followed by many accesses to device 2 followed by many accesses to device 1, etc. Each device could have been a bottleneck for a subperiod of the total monitoring period. This could also

DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

DEVICE UTILIZATION BY DEVICE NUMBER

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	DEVICE ID		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	REPORT 1
24	24	0.012	0.012	1-	1	I	
8736	8760	4.272	4.261	2-	2	Ixxxx	
24	8784	4.284	0.012	3-	3	I	
6584	15368	7.495	3.211	4-	4	Ixxx	
251	15619	7.618	0.122	5-	5	I	
10273	25892	12.628	5.010	6-	6	Ixxxxx	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
710	186175	90.802	0.346	46-	46	I	
36	186211	90.820	0.018	47-	47	I	
12	186223	90.826	0.006	48-	48	I	
25	186248	90.838	0.012	49-	49	I	
105	186353	90.889	0.051	50-	50	I	

Figure 8-16. Proportionate Device Utilization Report

```
                    DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                             DEVICE UTILIZATION BY DEVICE NUMBER

INDIV.   CUMUL.   CUMUL.   INDIV.   DEVICE            PERCENT OF OCCURRENCE                   REPORT
NUMBER   NUMBER   PRC      PRC      ID       00   10   20   30   40   50   60   70   80   90   100   2
                                             I....I....I....I....I....I....I....I....I....I....I
    9    186362   90.894   0.004    51-  51  I
    0    186362   90.894   0.       52-  52  I
    0    186362   90.894   0.       53-  53  I
   25    186387   90.906   0.012    54-  54  I
 2301    188688   90.028   1.122    55-  55  Ix
    .       .        .       .       .    .  .
    .       .        .       .       .    .  .
    .       .        .       .       .    .  .

205033 ENTRIES TOTAL               0 OUT OF RANGE
```

Figure 8-16.1. Device Utilization Report - Part 2

THIS PAGE LEFT INTENTIONALLY BLANK

have been the case if the proportionate utilization of each device was equal. The Channel Monitor can be used to uncover this cyclic type of usage. In addition, the Chronological Device Utilization Report (see subsection 7.5.18) was designed to uncover this type of problem by breaking down device utilization over time, rather than by utilizing a histogram. Nevertheless, when a single or small number of devices has a disproportionately large share of the accesses, they are potential bottlenecks and their usage should be further analyzed.

This report will show all connects that were issued to a given device. This includes all read/write connects, as well as any command type connects issued to a given device.(See subsection 7.3).

This histogram can report a maximum of 50 devices. If a site is configured with more than 50 devices, a second report will be produced, as a continuation of Report 1 (see figure 8-16.1).

8.5.9 Queue Length and Queue Time Histograms (File 57). Figure 8-17 shows the I/O queue length and queue time histograms for the I/O requests to devices as they are processed by IOS. These reports occur in pairs, one pair for each device. These reports are generated only for mass store devices. The first report in the pair is a length report and the second is a time report. The histogram is read in the same manner as the Proportionate Device Utilization Report. In addition to individual percentages, cumulative percentages are also reported. In figure 8-17 (part 2) we see that 90.279% of all requests have a queue time of 12 ms or less while 91.677% of all requests have a queue time of 22 ms or less. The number of entries in these two reports might not be equal. The first report is generated at the time of the request. When an I/O request is made, an entry is made to the histogram indicating the number of requests outstanding or in progress to this device. The second report is generated at the time of the actual connect and indicates the length of time between the request and the actual connect. Since observations have indicated that some STIOS are never connected, the first report may have a higher number of entries than the second. These two reports can be correlated by subtracting the number of STIOS not connected and the number of entries still in queue, for a given device, from the number of entries in the queue length histogram. Figure 8-17 (part 1) shows that 736 times there was a queue on device 9. This figure should correlate with the number of times device 9 was reported busy on the Channel Busy-Device Busy Report and the Channel Free-Device Busy Report. The two previous reports show the relationship between device and channel contention, while this histogram can be used for measuring the depth of queuing and the distribution of queue lengths.

Figure 8-18 shows the channel queue length and queue time for the I/O queue entries as they are used in the system channel environment. These histograms will be produced for both mass store and tape channels. Once again, the number of entries in these figures will not correlate because of STIOS not connected that were issued over a given channel, and not over a single device.

```
                          DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                        I-0 QUEUE LENGTH FOR IOM-0,PUB-08,DEVICE-09-- MS0450            DPC

INDIV.    CUMUL.    CUMUL.    INDIV.   QUEUE                  PERCENT OF OCCURRENCE                     REPORT
NUMBER    NUMBER      PRC       PRC    LENGTH   00   10   20   30   40   50   60   70   80   90   100    12
                                                 I....I....I....I....I....I....I....I....I....I....I
5848      5848      88.821    88.821    0- 0     Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
 639      6487      98.527     9.705    1- 1     Ixxxxx
  85      6572      99.818     1.291    2- 2     Ix
  12      6584     100.000     0.182    3- 3     I

6584 ENTRIES TOTAL      AVERAGE =   0.12834     VARIANCE =   0.149     STANDARD DEVIATION =   0.386
```

Figure 8-17. Device Queue Length and Time Histograms (Part 1 of 2)

```
                DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

              I-O QUEUE TIME-MS FOR IOM-0,PUB-08,DEVICE-09-- MS0450          DPC

INDIV.    CUMUL.    CUMUL.    INDIV.   QUEUE                 PERCENT OF OCCURRENCE                     REPORT
NUMBER    NUMBER     PRC       PRC      TIME     00   10   20   30   40   50   60   70   80   90   100    13
                                                  I....I....I....I....I....I....I....I....I....I....I
5828       5828    88.518    88.518     2.000    Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
  37       5865    89.080     0.562     4.000    I
  27       5892    89.490     0.410     6.000    I
  24       5916    89.854     0.365     8.000    I
  18       5934    90.128     0.273    10.000    I
  10       5944    90.279     0.152    12.000    I
  14       5958    90.492     0.213    14.000    I
  24       5982    90.857     0.365    16.000    I
  24       6006    91.221     0.365    18.000    I
  19       6025    91.510     0.289    20.000    I
  11       6036    91.677     0.167    22.000    I
  24       6060    92.041     0.365    24.000    I
 .          .        .         .         .       .
 .          .        .         .         .       .
 .          .        .         .         .       .
   2       6509    98.861     0.030    96.000    I
   3       6512    98.906     0.046    98.000    I
   2       6514    98.937     0.030   100.000    I

6584 ENTRIES TOTAL      AVERAGE =   6.32545    VARIANCE =   478.350   STANDARD DEVIATION =   21.871
  70 (1%) OUT OF RANGE    AVERAGE FOR THESE =   159.68    IN RANGE AVERAGE =   4.68
```

Figure 8-17. (Part 2 of 2)

```
                         DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                                    I-O QUEUE LENGTH FOR IOM-0,PUB-16

INDIV.    CUMUL.    CUMUL.    INDIV.   QUEUE                  PERCENT OF OCCURRENCE                       REPORT
NUMBER    NUMBER      PRC       PRC    LENGTH  00   05   10   15   20   25   30   35   40   45   50        219
                                                I....I....I....I....I....I....I....I....I....I....I
39161     39161     48.647    48.647    0-  0   Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
18926     58087     72.158    23.511    1-  1   Ixxxxxxxxxxxxxxxxxxxxxxx
12689     70776     87.920    15.763    2-  2   Ixxxxxxxxxxxxxxx
 6515     77291     96.014     8.093    3-  3   Ixxxxxxxx
 2349     79640     98.932     2.918    4-  4   Ixxx
  638     80278     99.724     0.793    5-  5   Ix
  146     80424     99.906     0.181    6-  6   I
   59     80483     99.979     0.073    7-  7   I
   11     80494     99.993     0.014    8-  8   I
    3     80497     99.996     0.004    9-  9   I
    2     80499     99.999     0.002   10- 10   I
    1     80500    100.000     0.001   11- 11   I

80500 ENTRIES TOTAL    AVERAGE =   0.96733    VARIANCE =   1.440    STANDARD DEVIATION =   1.200
```

Figure 8-18. Channel Queue Length and Queue Time Report (Part 1 of 2)

```
                DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                          I-O QUEUE TIME (MS) FOR IOM-0,PUB-16

INDIV.  CUMUL.   CUMUL.   INDIV.  QUEUE               PERCENT OF OCCURRENCE                REPORT
NUMBER  NUMBER   PRC      PRC     TIME     00   05   10   15   20   25   30   35   40   45   50     220
                                           I....I....I....I....I....I....I....I....I....I....I
22657   22657    28.182   28.132    2.000  Ixxxxxxxxxxxxxxxxxxxxxxxxxxxx
 2244   24901    30.973    2.791    4.000  Ixxx
 2289   27190    33.820    2.847    6.000  Ixxx
 2369   29559    36.767    2.947    8.000  Ixxx
 2433   31992    39.793    3.026   10.000  Ixxx
 2442   34434    42.830    3.037   12.000  Ixxx
 2381   36815    45.792    2.962   14.000  Ixxx
 2309   39124    48.664    2.872   16.000  Ixxx
 2106   41230    51.284    2.620   18.000  Ixxx
   .       .        .        .        .    .
   .       .        .        .        .    .
   .       .        .        .        .    .
  561   74099    92.168    0.698   92.000  Ix
  535   74634    92.833    0.665   94.000  Ix
  492   75126    93.445    0.612   96.000  Ix
  446   75572    94.000    0.555   98.000  Ix
  401   75973    94.498    0.499  100.000  I

80396 ENTRIES TOTAL     AVERAGE =    29.97953     VARIANCE =    1302.182     STANDARD DEVIATION =    36.086
  4423 (5%) OUT OF RANGE     AVERAGE FOR THESE =    125.58     IN RANGE AVERAGE =    24.41
```

Figure 8-18 (Part 2 of 2)

The user should refer to subsection 8.3 for a complete description of the methodology involved in making entries to these reports.

Figure 8-18 (part 1) shows that 41339 times there was a queue on channel 16. This figure should correlate with the number of times channel 16 was reported busy on the Channel Busy-Device Busy Report and the Channel Busy-Device Free Report. The two previous reports show the relationship between device and channel contention, while this histogram can be used for measuring the depth of queuing and the distribution of queue lengths.

8.5.10 Service Time Histograms (File 57). In all of the device histograms, it should be noted that the name of the device is also given. In figure 8-17, queuing statistics were presented for a device with the name DPC. If an exchange took place and the DPC disk pack was moved from 0-08-09 to 0-08-12, the data reduction program will account for that exchange and any connects that are made to 0-08-12 will be reported on this histogram and not to the 0-08-09 histogram.

In the upper right-hand corner of the report, a report number is indicated. This report number is used only to distinguish one histogram from another and in no way indicates the device to which the report refers. In addition, report numbers may not appear sequentially and this in no way is indicative of a problem. All histogram reports are generated by default and may be turned off by using the LIMITS option (subsection 8.6.13).

Figure 8-19 shows the I/O service time histograms for each I/O channel and device. Each histogram is given in 2ms intervals. The I/O service time is defined as the time (in ms) from connect to the time that IOS processes the terminate interrupt for the I/O request. These histograms are generated for both mass store and tape channels, as well as all mass store devices. On the bottom line of this report, an indication is given as to the percent of total time that this device or channel was busy.

The three device-oriented histograms, just described, have entries placed in them for every connect issued to the device (not just multicommands such as seek-read or seek-write).

```
                    DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                                I-O SERVICE TIME (MS) FOR IOM-0,PUB-16

INDIV.   CUMUL.   CUMUL.   INDIV.   SERIVCE             PERCENT OF OCCURRENCE                  REPORT
NUMBER   NUMBER    PRC      PRC      TIME    00   05   10   15   20   25   30   35   40   45   50     221
                                              I....I....I....I....I....I....I....I....I....I....I
  15        15     0.038    0.038    2.000    I
 390       405     1.012    0.975    4.000    Ix
1081      1486     3.715    2.702    6.000    Ixxx
1205      2691     6.727    3.013    8.000    Ixxx
1394      4085    10.212    3.485   10.000    Ixxx
1612      5697    14.243    4.030   12.000    Ixxxx
1806      7503    18.758    4.515   14.000    Ixxxxx
1895      9398    23.495    4.738   16.000    Ixxxxx
1978     11376    28.440    4.945   18.000    Ixxxxx
  .         .        .        .        .      .
  .         .        .        .        .      .
  .         .        .        .        .      .
 309     38337    95.842    0.773   92.000    Ix
 310     38647    96.617    0.775   94.000    Ix
 255     38902    97.255    0.638   96.000    Ix
 273     39175    97.937    0.683   98.000    Ix
 259     39434    98.585    0.648  100.000    Ix

40000 ENTRIES TOTAL     AVERAGE =   39.09831     VARIANCE =   667.607     STANDARD DEVIATION =   25.838
   566 (1%) OUT OF RANGE    AVERAGE FOR THESE =   104.80     IN RANGE AVERAGE =   38.16    % BUSY =   89.64
```

Figure 8-19. I/O Service Time Report (Part 1 of 2)

```
                  DISTRIBUTION COLLECTED ON SYSTEM AZ-DP3 AT 10:47:24 ON 83-08-04

                  I-O SERVICE TIME (MS) FOR IOM-0,PUB-08,DEVICE-09-- MS0450       DPC

INDIV.   CUMUL.    CUMUL.    INDIV.   SERIVCE            PERCENT OF OCCURRENCE                   REPORT
NUMBER   NUMBER     PRC       PRC      TIME    00   05   10   15   20   25   30   35   40   45   50     14
                                                I....I....I....I....I....I....I....I....I....I....I
  20       20      0.304     0.304     4.000   I
 136      156      2.370     2.066     6.000   Ixx
 227      383      5.819     3.449     8.000   Ixxx
 271      654      9.936     4.117    10.000   Ixxxx
 346     1000     15.193     5.257    12.000   Ixxxxx
 502     1502     22.820     7.627    14.000   Ixxxxxxxx
  .        .         .         .         .     .
  .        .         .         .         .     .
  .        .         .         .         .     .
  26     6203     94.242     0.395    90.000   I
  11     6214     94.409     0.167    92.000   I
  20     6234     94.713     0.304    94.000   I
  10     6244     94.865     0.152    96.000   I
  14     6258     95.077     0.213    98.000   I
  16     6274     95.321     0.243   100.000   I

6582 ENTRIES TOTAL     AVERAGE =    33.74268     VARIANCE =   1213.275    STANDARD DEVIATION =   34.832
 308 (4%) OUT OF RANGE     AVERAGE FOR THESE =   150.34     IN RANGE AVERAGE =   28.02    % BUSY =   12.73
```

Figure 8-19. (Part 2 of 2)

```
COL   1      8           16

      $      IDENT       1820251/30/3044
      $      SELECT      B29IDPXO/OBJECT/CM
      $      TAPE        01,X1D,,12345
      $      LIMITS      99,50K,-4K,30K
      $      DATA        I*

(The following is a suggested set of data cards)

JOB
2
TS1 FTS
LIMITS
SUMARY
END
```

Figure 8-25. CMDRP JCL

tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

8.9 <u>Tape Error Aborts</u>

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

```
********************************************************************************************************
********************************************************************************************************
*****                                                                                              *****
*****                          ***** T H E   C P U   A N D   T A P E   M O N I T O R *****        *****
*****                                                                                              *****
*****                                                                                              *****
*****                                                                                              *****
*****                            REDUCTION STARTED AT 1730:20.349 WED 82-12-27                    *****
*****              AND COMPLETED AT 1909:10.372 WED 82-12-27 A TOTAL TIME OF 1.65 HOURS           *****
*****                            ON SYSTEM NMCC2  RUNNING 4J2 ON TAPE D0020                       *****
*****                                                                                              *****
*****                       THE SYSTEMS CONFIGURATION CONSISTED OF:                                *****
*****                                                                                              *****
*****                                  4 - 6680  CENTRAL PROCESSORS                                *****
*****  PROCESSOR OPTIONS - 2K CACHE, EIS OPTIONS  INSTALLED, FAST MEMORY ACCESS, INSTRUCTION OVERLAP ON *****
*****                            2 PROCESSORS ARE ASSIGNED AND AVAILABLE                           *****
*****                                                                                              *****
*****                                       DISPATCHER OPTIONS                                     *****
*****                                  URGENCY THRUPUT SET                                         *****
*****                                  IN CORE PUSH AREA                                           *****
*****                                  DYNAMIC BUFFERING OF SSA MODS                               *****
*****                                  CLASS B PRIORITY ON                                         *****
*****                                  PRIORITY B JOBS                                             *****
*****                                                                                              *****
*****                              SNUMB          FREQ        TIME QUANTUM                         *****
*****                               TS1             2              7                               *****
*****                                                                                              *****
*****                                                                                              *****
*****                 MONITOR OVERHEAD  -  MONITOR     TIME(SEC)     % OVERHEAD                    *****
*****                                      -------     ---------     ----------                    *****
*****                                       EXEC           47            3.26                      *****
*****                                       CPU            47            3.26                      *****
*****                                      -------     ---------     ----------                    *****
*****                                       TOTAL          94            6.52                      *****
********************************************************************************************************
********************************************************************************************************
```

Figure 11-2. CPU Title Page

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

            DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

                               CPU QUEUE LENGTH DISTRIBUTION

 INDIV.   CUMUL.   CUMUL.   INDIV.   QUEUE        PERCENT OF PROBABILITY OF OCCURRENCE        REPORT
 NUMBER   NUMBER    PRC      PRC     LENGTH   00   10   20   30   40   50   60   70   80   90   100     1
                                               I....I....I....I....I....I....I....I....I....I....I
   231      231    84.00    84.00     0-   0   IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
    44      275   100.00    16.00     1-   1   IXXXXXXXX

  275 ENTRIES TOTAL    AVERAGE =    0.16000    VARIANCE =    0.134    STANDARD DEVIATION =    0.367
```

Figure 11-3. CPU Queue Length Distribution

```
$     ALTER   8258                                                      ***GMF**
      TRMAC   19          TRACE SNUMB ISSUING WRITE                     ***GMF**
$     ALTER   8300                                                      ***GMF**
      IFE     GMF,1,1                                                   ***GMF**
      SYMDEF  TP3115                                                    ***GMF**
$     ALTER   8338                                                      ***GMF**
      TRMAC   20          TRACE SNUMB ISSUING READ                      ***GMF**
$     ALTER   8624                                                      ***GMF**
      TRMAC   5           TRACE # AVAILABLE UNDOT ENTRIES               ***GMF**
$     ALTER   9278                                                      ***GMF**
      TRMAC   21          TRACE TPAP ABORTED                            ***GMF**
$     ALTER   10707
   (NTPE ALTER cards)
      TRMAC   22          TRACE TPAP BOMBED, WHY                        ***GMF**
   (NTPE ALTER cards)
```

Figure 12-21. (Part 11 of 11)

```
$           CHANGE        19,19
GMF         SET           1

$           CHANGE        1223,1224
            LDX2          TP0316
            SXL2          TRACE1+2

$           CHANGE        1229,1230

$           CHANGE        1232,1233
            STA           TRACE1+2

$           CHANGE        1237,1238

$           CHANGE        1241,1242
            STA           TRACE1+2

$           CHANGE        1262,1262
            LDA           SREGS+5

$           CHANGE        1264,1264

$           CHANGE        1266,1266

$           CHANGE        1326,1326
            TRA           T74.21

$           CHANGE        1364,1364
            TRA           T74.21
```

Figure 12-22. Transaction Processing System CHANGE Cards (Part 1 of 3)

```
$           CHANGE      1725
            SYMREF      TREP
GMF         SET         1
TRMAC       MACRO
            IFE         GMF,1,2
            XED         TREP
            VFD         18/#1,18/60 0-17=SUB.TYP,18-35=TYP.74
            ENDM        TRMAC

$           CHANGE      1757
            IFE         GMF,1,1
            TRMAC       4

$           CHANGE      1847
            IFE         GMF,1,1
            TRMAC       4

$           CHANGE      2282,2282
GMF         SET         1

$           CHANGE      3086,3089
            IFE         GMF,1,2
            LDQ         MBUF+.MIT
            TRMAC       35
            MLR

$           CHANGE      3444
            IFE         GMF,1,2
            LDQ         .MID
            TRMAC       35
```

Figure 12-22. (Part 2 of 3)

$	CHANGE	3963,3963
$	ALTER	5289
$	CHANGE	4008,4008
$	ALTER	5729
$	CHANGE	4016,4016
$	ALTER	5808
$	CHANGE	4805,4805
GMF	SET	1

Figure 12-22. (Part 3 of 3)

SECTION 14. CONDUCTING A SITE COMPUTER PERFORMANCE EVALUATION USING THE GMC

14.1 Introduction

A general plan for using the GMC System and document when conducting a performance evaluation of a Honeywell 6000 computer system is presented in this section. Detailed analysis procedures that guide the analyst in applying specific techniques to analyze system performance are introduced. The primary purpose is to aid the analyst in his use of the GMC reports. This section will indicate which reports the analyst should reference, what data he should extract from those reports, and some guidance as to what the reports are indicating. In some cases, possible corrective ("tuning") recommendations will be made, but tuning guidance is not the primary purpose behind this section. If a tuning study should be made, the user must rely on his own personnel skills and the tools of GMF to perform such a study.

14.2 General Definitions

14.2.1 Computer Performance Evaluation (CPE). Computer performance evaluation is a generic term applied to many techniques for determining the performance characteristics of both a computer system and its associated site processing activities. The performance characterisitics may be compared to many criteria, including: (1) standards of economic operation, (2) technical norms, or (3) measures of service provided.

14.2.2 Computer System Performance Variables. The performance of a computer system is influenced by nearly every facet of the data processing function. The following examples define the scope of the computer performance function.

14.2.2.1 System Design. Computer application system design and development can be the starting point of performance degradation. Errors in original design with respect to I/O media selection, file structures, frequency of run, etc., may result in less than optimal performance for as long as an application is in existance.

14.2.2.2 Programming. A computer programmer's proficiency and the availability of program optimization tools, for example, will influence program design and coding, and affect system performance.

14.2.2.3 Hardware Configuration. Specific components of a computer system may be mismatched to the system as a whole, causing major subsystems (or the entire configuration) to operate at a reduced performance level. Even if the performance capabilities of the individual subsystems are reasonably well-matched, the system may be poorly configured for the site's workload, resulting in poor performance.

14.2.2.4 System Software. System software is considered to be those programs supplied by the mainframe vendor. This software may be inappropriately parameterized to fit the site workload, or may be a source of high overhead or bottlenecks to efficient workload processing.

14.2.2.5 Operations. Operations consists of functions such as the scheduling of the workload, providing job assembly (and library) services, and operating the system through the console. All of these functions are vital to the proper operation of the system. Mistakes, insufficient training, poor documentation, and a variety of other factors may contribute to operational problems which substantially decrease system performance.

14.2.2.6 Communications Hardware and Software. A communications network, its interface to a central system, and the software used to control the on-line applications, may have a significant impact on the system's overall performance.

14.2.2.7 Computer System Performance Tuning. The process of analyzing and appropriately adjusting computer system performance variables is known as computer system performance tuning. After becoming proficient in the use of GMF, the user should be able to conduct his own tuning study.

14.2.2.8 Turnaround Time. This is the total elapsed time taken by a job (or set of jobs) submitted to a WWMCCS site for batch processing. Total batch turnaround time is comprised of computer system processing and physical input and output handling in the machine room, both before and after system processing. A job's total turnaround time therefore includes all processing and waiting points through which the job must pass from submission until return to a user. The GMC measures only a batch job's computer processing turnaround time. The other measure of total turnaround time (input/output handling) are important and can be a bottleneck to total turnaround time.

14.2.2.9 CPE User Objectives. The definitions associated with any or all of the above data processing functions can be considered variables that affect computer system performance. The GMC reports should be used to determine which of the above functions need further investigation.

14.3 Solutions to Performance Problems

Particular resource bottlenecks may be confirmed as elongating turnaround or response time. Several solutions can usually be applied to remove a particular bottleneck. In general, four kinds of solutions exist to remove identified bottlenecks.

14.3.1 Scheduling Solutions. These solutions change the way that either batch or TSS workloads are scheduled for processing. They shift particular workloads to more evenly distribute system resources across the workload.

14.3.2 Parameter Solutions. These changes involve adjustments to system or subsystem functions. Examples include: (1) changes to the parameters of the GCOS Dispatcher or (2) a change in the placement of GCOS libraries. A solution may also include specific changes to GCOS code, made through authorized scftware patch procedures.

14.3.3 Programming Solutions. These changes can involve modification of one or more processing jobs running in the system. Recommendations may be made to speed application of jobs by changing particular file locations discovered as delaying the job.

14.3.4 Sizing Solutions. These types of system changes involve an increase, decrease, or realignment in the system's hardware configuration.

14.4 Structure of the Analysis Process

The analysis process in figure 14-1 is a flow chart that should be referenced during a performance study. The analysis process is comprised of two phases: (1) a Problem Definition Phase and (2) a Problem Analysis Phase. The activities of the Problem Definition Phase are directed toward determining whether a batch turnaround time or TSS response time problem actually exists. The activities of the Problem Analysis Phase are directed toward revealing causes of the identified turnaround time or response time problem.

14.4.1 Starting The Process. The evaluation process may begin in one of three ways: (1) by direct request, (2) by an internal review of site-selected service elongation metrics, (3) by user requests.

14.4.1.1 Direct Request. The request may by initiated by WWMCCS management, directing a facility to perform an analysis and tuning effort. Many reasons could cause management to request a CPE study. Examples include: (1) desire to extend the life of a system, (2) pressure applied from users through higher authority, (3) awareness of the potential for performance gain after an evaluation, and (4) a desire for management to evaluate cost reduction programs.

14.4.1.2 Internal Review. The decision to initiate a study may result from the output of a performance exception reporting system or from the desire to conduct a periodic review of site operations.

14.4.1.3 User Requests. Users of site services may request an evaluation. Complaints of unacceptable batch turnaround time or TSS response time can initiate a search for their causes. Unfavorable comparisons with service rates provided at other installations can point out the need for a study to determine if service can be improved.

14.4.2 Problem Definition Phase. The Problem Definition Phase (see figure 14-1, part 1) is comprised of four activities: (1) define and verify the

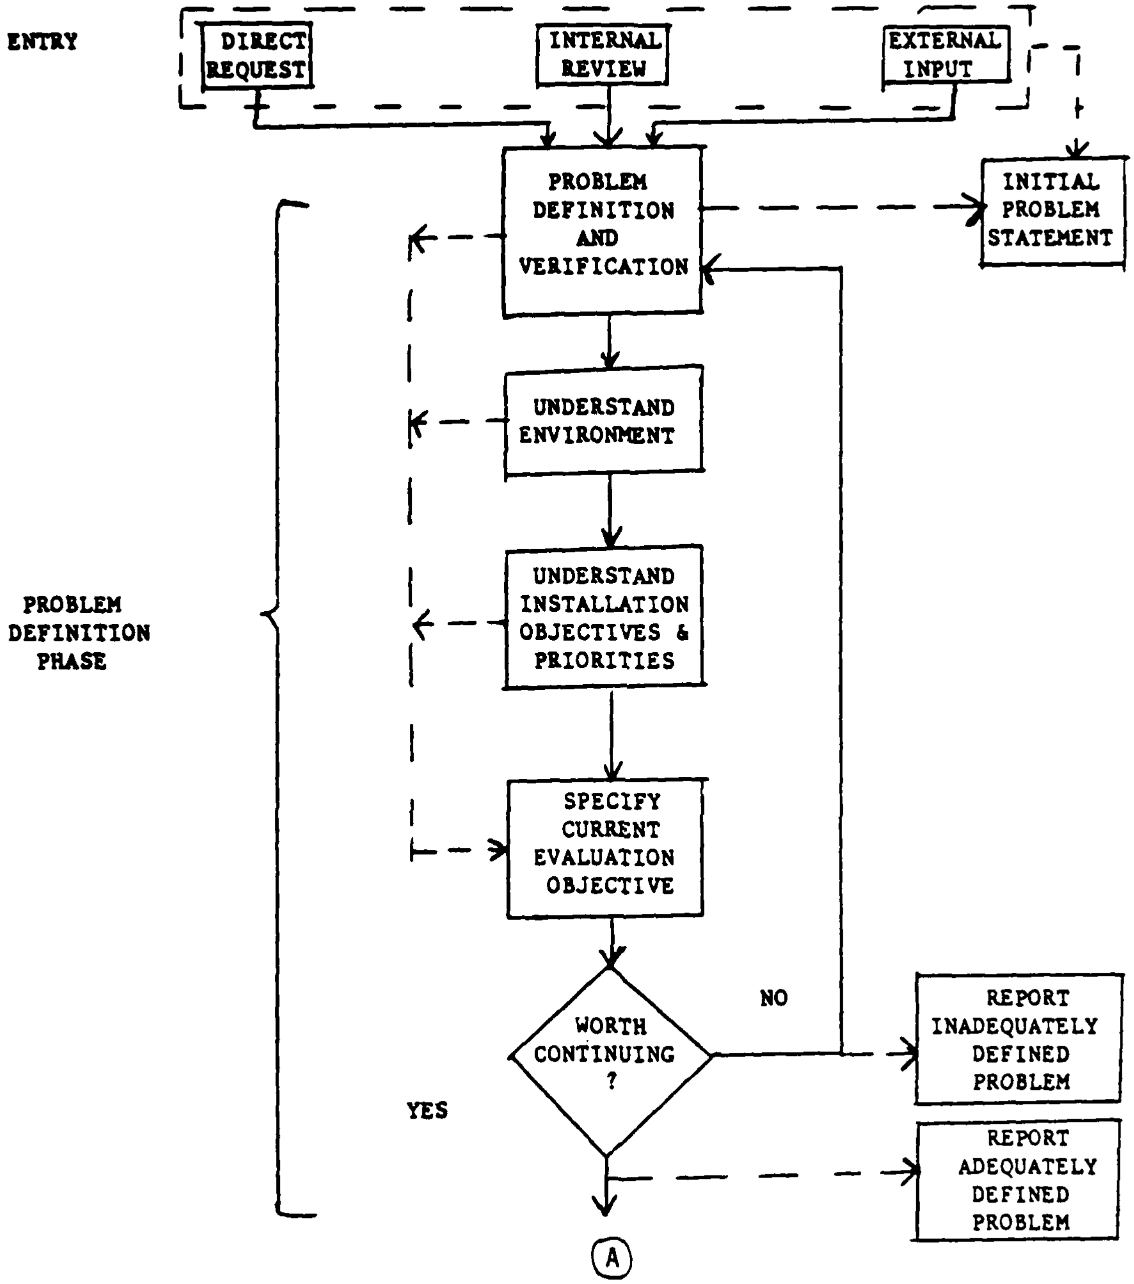

Figure 14-1. Flowchart Of Evaluation/Tuning Process (Part 1 of 2)

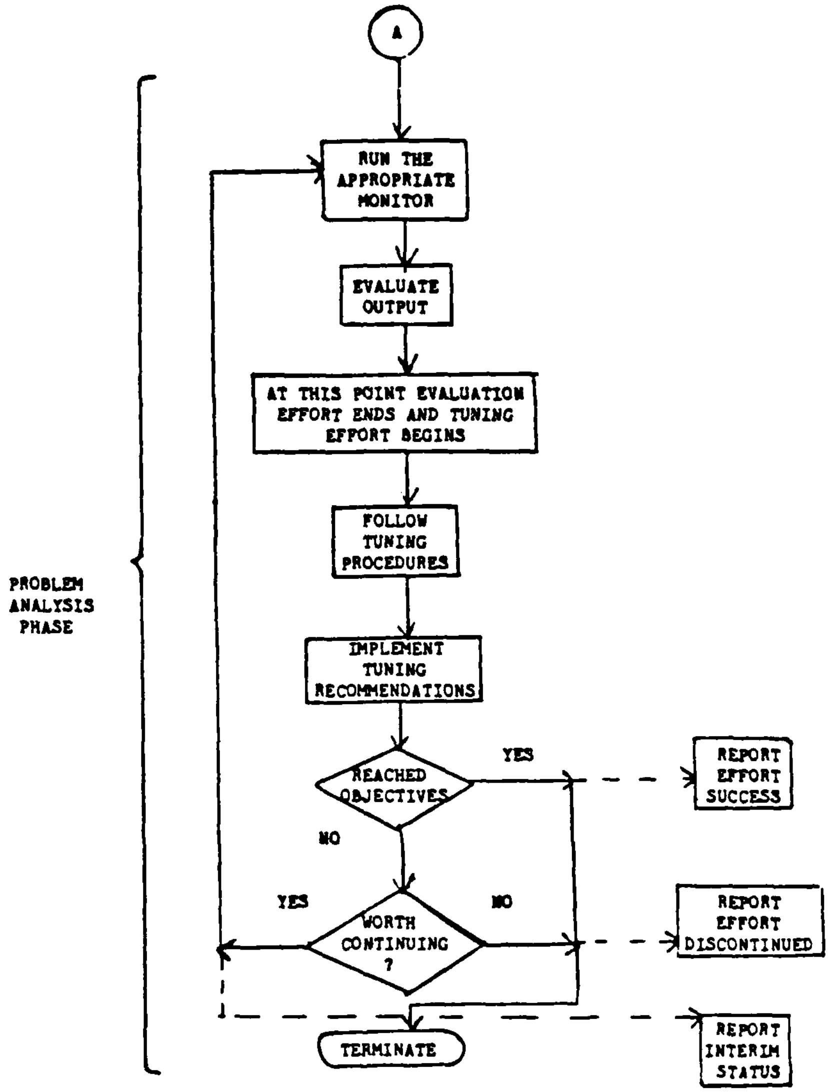

Figure 14-1. (Part 2 of 2)

problem, (2) gain understanding of facility environment, (3) understand installation service objectives and priorities, and (4) specify current evaluation objectives. Some of these activities are initiated as evaluation begins. Others are maintained as on-going activities. The activities of the Problem Definition Phase are introduced below.

14.4.2.1 Define and Verify the Problem. For whatever reason the evaluation was initiated, a starting point or premise must be established. The problem must be described as well as it can at this point, even if at a very general level. "Batch turnaround time is poor" or "TSS response time must be improved" are valid statements of problems at this time. Site management must then verify that this is the problem they wish to pursue. The facility staff and users should verify that, in fact, this problem does exist and that they view it as a problem of importance. Note that a specific evaluation objective is not yet defined; only the basic problem statement is verified.

14.4.2.2 Gain Understanding of Facility Environment. This second step helps the CPE analyst relate the system environment to the problem statement. This activity may be time-consuming if the analyst has little personal experience with the facility and the system. Information collected during this activity will help determine the reason for the problem's importance; it may explain a reason for the problem's existence and help rank one problem against another. The activity forces an analyst to view the entire facility; this is important since many performance problems are caused by combinations of factors. The activity assists the analyst in understanding how to narrow, refine, modify, and improve the problem's definition in order to attempt a valid solution. Each area described in the following paragraphs should be examined. It may help the analyst to examine them in the order they are given below.

14.4.2.2.1 Hardware Configuration. The exact system configuration should be determined and a diagram of the overall structure of the system should be created. Collection of reliability statistics on major configuration components and a history of hardware changes to the configuration should be done.

14.4.2.2.2 Software Configuration and Development Practices. The exact operating system configuration should be determined. Any local changes made to the system and any specialized major subsystems that are running should be documented. Document software monitors and other CPE measurement techniques used. Determine program optimization techniques used and note standards imposed on operations and programming staffs.

14.4.2.2.3 Existing CPE Practices. Responsibility for computer performance evaluation at the site must be determined. Identify sources of CPE data that are employed and determine how the data are used to evaluate system performance. Determine if personnel in all site functional areas (e.g., programming and operations) can relate to the performance data. Research how previous CPE studies conducted at the site in the past have been documented.

14.4.2.2.4 Site Workload Characteristics. The existing workload uses of system resources should be described. Identify patterns of resource use by selected jobs. Obvious bottlenecks to the handling of jobs within the installation should be determined. The resource monitor, the RMC portion of GMF, should prove very useful during this phase.

14.4.2.2.5 System Users. The workload analysis should be expanded by investigating the practices of the major users in the installation. Determine how special priorities are assigned to these users and whether the users can directly control system resources. Note chargeback schemes used (if any) to level the workload across the operating day and night. Evaluate the levels of user satisfaction (or dissatisfaction) exhibited.

14.4.2.2.6 Operations Practices. The operating shift schedule for the site should be examined. Analyze pre-scheduled and nonproduction time and unscheduled nonproduction time. Evaluate training provided for operators and systems programmers. Study site library maintenance procedures. Examine the production control points established by operations. Observe how formalized logs are maintained to track work as it passes through the installation.

14.4.2.2.7 Batch Job Scheduling. Observe how batch work is accepted for processing at the site. Study the techniques used to schedule the batch workload that are either automated or are manually implemented.

14.4.2.2.8 Site and Computer Facility Organization. Study the organization structure and reporting authority at the site. This includes the organization of the site itself. Determine the extent to which system sizing activities are organizationally separated from operations and applications systems development.

14.4.2.2.9 CPE Checklist. The use of a CPE checklist will greatly aid the analyst in his understanding of the facility environment. Figure 14-2 is an example of such a checklist.

14.4.2.3 Understand Installation Service Objectives and Priorities. This third step determines site processing objectives and priorities to be applied to decisions made during both the Problem Definition and Problem Analysis Phases.

14.4.2.3.1 Installation Service Objectives. To be effective, a CPE effort must consider the hardware, software, personnel, and service objectives of the computer installation being analyzed. Although the first three areas are readily identifiable, the site's service objectives are often misunderstood. An installation's objectives can emphasize production or

A. Configuration:

____1. MPC/PSIA crossbar check

____2. Trace setting/SCF collection

____3. CATDUP off unless device saves/restores are done

____4. NIAST buffer count

____5. TSS swap file location/size (2200)

____6. Placement of system files on low-use devices

____7. SYSOUT files 4 or 6 -- do not use 5

____8. No module CKSUM

____9. SSLOAD maximum 20-30 based on workload

____10. IOM-lines

____11. HCM available for other modules

____12. SCOMM - set for intercom I/O area. Should not be 128

B. Options:

____1. Dispatcher

  ____a. Urgency thruput

  ____b. MCOUNT off

  ____c. Priority B (why)

____2. TSS

  ____a. Devices for swap files

  ____b. Large user penalty

  ____c. Priority B (NO)

  ____d. Round Robin swap files (Reston patch number)

  ____e. 4 swap files - large system

Figure 14-2. CPE Checklist (Part 1 of 3)

____3. FMS Cache #A2E

____4. Memory move quantity #AOG

____5. SSA in memory pushdown #A1E
bit 3 = 1

____6. SSA Cache size #A1Y

____7. GRTS

____a. DBIAS,BBIAS = 1000_8

____b. MAXMSG = 6 if load 5600_8

____c. Line size setting

____d. Trace setting

C. Hardware:

____1. Memory interleaving

____2. MPC/PSIA connection

D. Miscellaneous:

____1. MSCAN to set up scheduler classes based on required resources

____2. Limit use of high urgency

____3. USERID hash

____4. Cold Boot frequency (FRAG)

____5. TPAPs not needed

____6. Jobs waiting peripherals on VIDEO being reviewed by operator
(overdue allocation problem)

E. Memory Evaluation

F. CPU Evaluation

Figure 14-2. (Part 2 of 3)

G. <u>Communication Evaluation</u>

H. <u>I/O Evaluation</u>

Figure 14-2. (Part 3 of 3)

availability, or a mixture of both. A mixture is generally dominated by either a production or availability objective. A production objective attempts to take full advantage of the capabilities of the system in terms of throughput. With a customer-oriented availability objective, however, the interests of the users take priority over the most efficient use of system resources. In either case, the computer installation's management objective is return on investment. For the production objective, the most important investment is the computer; for the availability objective, the most important investment is the people or system using the computer. The difference between these two objectives is reflected in how the performance parameters are evaluated.

14.4.2.3.1.1 Production Objective. Traditional data processing is production-oriented. The computer installation is viewed as a large investment for capacity to do routine work. Much of that capacity depends on a high degree of simultaneous use of many system components. Management's production objective is to use as many of the components as much of the time as possible by scheduling compatible jobs. Under these conditions, low batch turnaround and TSS response times are secondary to high machine utilization. Management's success is generally defined in terms of a relatively high number of jobs contending for computer resources (a high multiprogramming level), high central processor unit (CPU) activity, nearly full memory utilization, and highly active channels. Schedule-driven processing uses internal and external priority allocations to sustain efficiency. Programs must be written to share available resources. This can be done with program segmentation techniques, use of a minimum number of devices, the indirect use of input/output through spooling, and adherence to rigid standards -- all at the expense of the individual job. Data must be so distributed that device activity is economically justified with the attainment of a low system wait, low unit cost, and high simultaneity.

14.4.2.3.1.2 Availability Objective. Service-oriented systems are more concerned with return on investment for users (managers, scientists, or programmers) than with the computer itself. Low TSS response and batch turnaround times are critical. Their increase delays user operations on such applications as on-line command and control systems, real-time update systems, scientific support services, and program development systems. Demand-driven processing varies in activity levels, both by users and type of work. Minimum emphasis is placed on scheduing. The system must therefore have available capacity ready to respond to demand. Such a system must have utilization rates well below the 100% limit in order to accommodate the variations in demand. Success is measured in terms of user satisfaction and little emphasis is placed on reporting high utilization figures.

14.4.2.3.1.3 Mixed Objectives. Usually computer installations encounter (1) demands for highly responsive services and (2) pressure from management for high production rates. The two objectives are not mutually exclusive.

Predominance of either objective can be identified within operational time periods, and management must evaluate whether the satisfaction of one objective might deleteriously affect the satisfaction of another.

Computer performance evaluation can serve the management of either type of computer installation and can also serve mixtures that might have different objectives, depending on the time of day or day of the week. However, it is important that managers, auditors, and executives recognize the implications of the different objectives when they compare the performance of one installation to another.

14.4.2.3.2 Installation Priorities. The installation's priorities derive from the predominance of either the production objective or the availability objective at a particular site. They also, by implication, determine the sequence in which system tuning solutions are applied.

14.4.2.3.2.1 Service Priorities. A site may feel that low TSS response time is more desirable than low batch turnaround time for a certain period of the operating day. The analysis procedures presented in this section are not generally directed toward determining which of the two is "more important". However, the sequence of examining either of the two service areas may be affected by this priority.

14.4.2.3.2.2 Evaluation/Tuning Solution Priorities. Tuning solutions are presented to correct system bottleneck conditions. In nearly all cases, more than one solution is specified to correct a problem. The solutions are generally proposed in a sequence that recommends the more quickly (or easily) applied solution be implemented before others. Installation requirements may change this sequence.

14.4.2.4 Specify Current Evaluation Objective. This fourth step is used to determine whether the originally specified problem can actually be investigated with currently-available procedures.

14.4.2.4.1 Objective. An objective is a stated (i.e., documented) goal of an analysis. An objective must be stated in specific and quantified terms. The objective statement is used to determine when a particular analysis has been completed. Examples of well-stated objectives are: (1) "reduce mean response time for (stated) non-trivial TSS commands to five seconds" and (2) "reduce the mean batch turnaround time for (stated) jobs to 1.5 hours." A well-stated objective includes: (1) definition of the workload catagory, (2) a description of the process to be investigated, and (3) a service metric value. Examples of badly-stated objectives are "reduce TSS response time" and "improve turnaround time." Note the missing components of these objectives.

14.4.2.4.2 Objective Decision. Determine whether the objective as documented is realistic, attainable, and a cost effective target.

14.4.2.4.2.1 Attainable Objectives. Objectives must be attainable. It might be possible to reduce a program's elapsed time to certain limits, but not to a desired limit, simply because of the quantity of I/O and computation that occurs within the program. If this is the case, re-evaluate the objective.

14.4.2.4.2.2 Realistic Objectives. It might be possible to reduce a program's elapsed time to the desired amount. However, particular site constraints may prevent its being a realistic goal. These constraints might have their source outside of the tuning effort, but directly affect the internal performance of the program. These might include, for example, the requirement to give certain other jobs higher priority. If this is the case, re-evaluate the objective.

14.4.2.4.2.3 Cost-Effective Objectives. The additional effort, cost, and time required to achieve an attainable and realistic objective might not be worth it. A particular increment of performance improvement might simply not be cost effective. Determine the amount of improvement likely with additional effort. Decide whether the effort might be better off abandoned.

14.4.2.4.3 Determine if Worth Continuing. The final part of step four is to determine whether the process itself is worth continuing. There may be potential performance problems of greater magnitude or immediate importance that may have been uncovered during the Problem Definition Phase. Document the decision and the objective and obtain management approval of them.

14.4.2.4.4 Begin Problem Analysis. Having obtained concurrence from management, begin the second half of the performance evaluation process (see figure 14-1, part 2): the Problem Analysis Phase. Activities of this phase involve the execution of the various GMC monitors and the analysis of their output.

14.4.3 Problem Analysis Phase. The Problem Analysis Phase (see figure 14-1, part 2) is comprised of five activities briefly described below:

14.4.3.1 Run Appropriate Analysis Tool. Analysis requires collection of data to start an investigation. The various GMC monitors have been developed for this purpose.

14.4.3.2 Evaluate System Output. The output from the various GMC monitors must be analyzed.

14.4.3.3 Follow Tuning Procedures. The user must develop his own skill in developing provide specific procedures to test hypotheses by examining the reports produced by the GMC monitors. Specific system tuning steps can result from the tests.

14.4.3.4 Evaluate Need to Continue the Tuning. This decision is made after the relevant recommendations have been implemented. If the current objective (specified in the last activity of the Problem Definition Phase) has not been met, the user analysis can define areas for further investigation.

14.5 Composition of a Performance Evaluation Team

While this section has been developed to aid WWMCCS ADP Managers in the application of various computer performance evaluation tools and techniques, the actual team conducting this analysis must have a reasonable education and experience level on the H6000. The team should be comprised of at least one individual familiar with system software and its operation; one individual generally familiar with system hardware (at least in terms of functionality); one individual familiar with the procedures for executing jobs on the H6000 and at least one individual (probably a manager) who is familiar with the objectives and priorities of the installation. It certainly is possible that one individual may be able to handle several of the above functions.

14.6 System Evaluation

14.6.1 Introduction. A total system evaluation should be performed at least once a year. It may have to be performed more often, depending upon the results of daily resource statistics. A steady increase in resource utilization would possible imply the need for an evaluation. The Resource Monitor (RMON) of GMF should be used to track daily resource utilization.

The data collected for the detailed system evaluation must be representative (i.e., typical) of the data that would be collected on a "normal" day. This requires the collection of data over several time periods. The suggested schedule for collecting data for a total system analysis is to run the GMC monitors intermittently for two weeks. The monitors should be run as two separate groups, on alternate days. Group one would consist of the Memory Utilization Monitor (MUM) (possibly the Idle Monitor (IDLEM)), CPU Monitor (CPUM), Communications Analysis Monitor (CAM), and the GRTS Monitor (GRTM). Group two would consist of the Mass Storage Monitor (MSM), Channel Monitor (CM), and possibly the Idle Monitor (IDLEM). This sequence provides a good representative sample of the system workoad. To limit the amount of data collected, it is advisable to run Group Two only during prime time, heavy usage (4-6 hours per day). It may also be desirable to run the combined set of monitors at least once per week during the two-week monitoring session. This allows the analyst to get a complete, unified picture of the entire performance of the system. It must be realized that a large number of tapes can be generated under such a monitor configuration. Therefore, data collection should be limited to no more than four hours (during heavy prime time usage).

14.6.2 Selecting a Representative Value From GMC Histogram Reports. Some test procedures require an analyst to pick a "Representative Value" to describe a frequency distribution. The following paragraph gives guidelines for choosing Representative Values based on the type of distribution observed.

Figure 14-3 shows a hypothetical "Total Elapsed Time an Activity Was in Memory" report. Variations of this report will be used to illustrate types of distributions. This chapter will reference only the pictorial part of the report. All other parts of the histogram report are described in Chapter 6 of this document.

14.6.2.1 Symmetric Distribution Closely Clustered Around a Single Point. The "Representative Value" for the distribution shown in Figure 14-3 is easy to select. The values are clustered around the "Average": 17.6 tenths of a second, or 1.76 seconds. The shape of the distribution is symmetric -- about the same number of activities had values over 1.76 seconds as under 1.76 seconds, and the standard deviation is small when compared to the average. The absence of a second line under the "Entries Total" line indicates that no activities stayed in memory longer than 2.4 seconds. When a distribution resembles this example, use the "Average" printed at the bottom as the Representative Value.

14.6.2.2 Skewed Distribution. Figure 14-4 shows another distribution. This distribution is "skewed" (i.e., not symmetric), because most of the activities spent around 0.1 to 0.7 seconds in memory, while some spent as much as four or five seconds. Care should be taken when selecting a "Representative Value" from this distribution. If the analyst wants to emphasize the "typical" activity, which stayed in memory 0.3 seconds or less, he could select the "median" (the value which evenly divides the activities in the distribution -- half spent less time in memory, and half spent greater time in memory). In figure 14-4, the median is about 0.29 seconds. The median can be estimated from these reports by descending down the "Cumulative" column (not displayed in the figure) until the value first exceeds 50. The median falls within the time range of this row.

14.6.2.3 Distribution With Outliers. There are some instances when the distribution will also have some values that were too big to fit in the histogram. This condition will be indicated by an additional output line at the bottom of the report. This line will indicate the number of occurrences that were outliers, the average for just the outliers, and the average for the values that fit into the report, minus the outliers. The three important factors about these types of distributions are: (1) the amount of times they occur; (2) the percent of the total values that are outliers; (3) the "in-range average."

If the percent of outliers is greater than 10%, the analyst should use the overall average, given in the first line of the report, for any comparisons. If the percentage of out-of-range values is less than 10%, the analyst can use the "in-range average" value for his comparisons since the effect of the outliers will most likely be minimal on the total system performance.

DISTRIBUTION COLLECTED ON SYSTEM NMCC at 14:41:24 on 81-05-15

The Total Elapsed Time An Activity Was In Memory

```
                      Percent of Occurance           Report 21

Tenths   00   05   10   15   20   25   30   35   40   45   50
Second    /    /    /    /    /    /    /    /    /    /    /
10-10     I
11-11     I
12-12     IX
13-13     IXX
14-14     IXXXX
15-15     IXXXXXXXXXX
16-16     IXXXXXXXXXXXXXXX
17-17     IXXXXXXXXXXXXXXXXXX
18-18     IXXXXXXXXXXXXXXXXXX
19-19     IXXXXXXXXXXXXXXX
20-20     IXXXXXXXXX
21-21     IXXXXX
22-22     IXX
23-23     IX
24-24     I

9276 Entries Total   Average 17.62   Variance 4.85   Standard Deviation 2.20
```

Figure 14-3. Sample Symmetric Distribution

```
DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 14:14:24 ON 81-05-15

       The Total Elapsed Time An Activity Was In Memory

                    Percent of Occurence

     Tenth    00   05   10   15   20   25   30   35   40   45   50
     Second    /    /    /    /    /    /    /    /    /    /    /

      0-1      Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
      2-3      Ixxxxxxxxxxxxxxxxxxxxxxxxxxx
      4-5      Ixxxxxxxxxxxxxxxx
      6-7      Ixxxxxxxxxxx
      8-9      Ixxxxxx
     10-11     Ixxxx
     12-13     Ixxx
     14-15     Ixx
     16-17     Ix
     18-19     Ix
     20-21     I
       .
       .
       .
     56-57     I

     118669 Entries Total  Average 4.247  Variance 18.52  Standard Deviation 4.3
```

Figure 14-4. Sample Skewed Distribution

14.6.3 Memory Evaluation. The first step in a memory evaluation is to summarize all the pertinent information. The easiest way to do this is for the user to create a table similar to figure 14-5. The information for each column is collected from various MUM reports. Once the chart is filled out, the analyst can then ascertain a reasonable idea of the overall memory status of the system. The reports required to collect the statistics will be discussed, as well as an analysis procedure. The current version of the MUM will automatically produce this table (see subsection 6.3.13 - Memory Statistics Report).

14.6.3.1 Obtaining the Data. The first step is for the user to examine the monitor data collection reports for the total time of the run. Any monitoring session of less than 2 hours in duration should be discarded. The Title Page will indicate the overhead generated by each of the GMF monitors active during the data collection phase. While this information is not used in this analysis, it is an item of information that is usually requested. Columns 1 and 2 of figure 14-5 can be obtained directly from the MUM Title Page. Under the System Configuration section of the Title Page, the amount of memory configured on the system will appear. This figure should be tentatively noted under column 5. Following this figure, the Title Page will list the amount of memory used by Hard Core, Core Allocator, SSA Cache and if any memory releases occurred during the monitoring session. All these functions have the effect of reducing the available user memory. These values should be summed and tentatively recorded under column 6.

Following the above information are several lines of statistics concerning the processing characteristics of the system. Columns 17 and 18 may be filled from this information. When determining the number of activities processed per hour, the analyst has two figures available. The analyst may choose to record the total number of activities processed per hour or he may record the number of (non-system schedular) activities processed per hour. During the course of the day many system schedular activities (activity 0 of any batch job) are processed. These activities are not really user generated, but rather are generated from the system. Therefore, by removing these activities from the total activities processed, a more realistic figure will be generated. The final three lines of the title page can be used to provide the data for columns 3 and 7 on the chart. Column 4 is filled from Report 1, columns 8 and 9 from Reports 11 and 12, columns 10 and 11 from Reports 16 and 17, column 12 from Report 19, column 13 from Report 51, column 14 from Report 37, column 15 from Report 50, and column 16 from Report 36. The two remaining columns that need to be completed are columns 5 and 6. The System Program Usage of Memory Report should be used to complete column 6. When processing the MUM data reduction program, the user should seriously consider using the MASTER input option. This will provide the user with a much better indication of the true system program load. In order to complete column 6, the user should record the total value appearing under the "WEIGHTED TTM" column of the System Program Usage of Memory Report. This value should then be added to the value already recorded under

#1	#2	#3	#4	#5	#6	#7
Date	Hours	CPU/IO Ratio	Avg Activity Size (K)	Amt of Memory available for User activities	Amt of Memory used by System Functions	Memory Surplus or Memory Shortfall with and without the PALC Queues

#8	#9	#10	#11	#12
Avg Number of User Activities Waiting Memory	Avg Number of System Activities Waiting Memory	Avg Number of User Activities in Memory	Avg Number of System Activities in Memory	Ratio of Duration vs Memory Time

#13	#14	#15	#16	#17	#18
% Slave Memory Used	Amount of Time User Activity Swapped	Amount of Time Activity Waited for Original Memory Allocation	Number of User Swaps	Activities Per Hour (Throughput)	Total Swaps Per Hour

Figure 14-5. Memory Statistics

AD-A138 533 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 7 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 01 OCT 83 CCTC-CSM-UM-246-82-CHG-7

3/3

UNCLASSIFIED F/G 9/2 NL

END
DATE
FILMED
4-84
DTIC

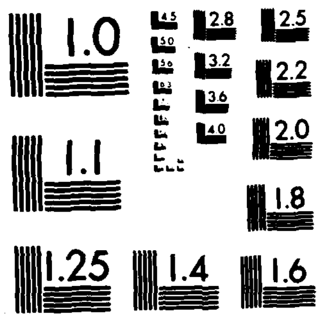

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

column 6. Finally column 5 can be determined by subtracting the value reported in Column 6 from the value previously reported in column 5.

14.6.3.2 Evaluating the Data. Figure 14-5 should be used in the following manner to determine if memory is a system constraint.

Step 1 - If column 7 shows a surplus of memory greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that there is no memory constraint on the system. If Column 7 shows a surplus of memory, but does not exceed the aforementioned limits, the implication is that the current system has sufficient memory, but that the system is approaching memory saturation. If Colulmn 7 shows a shortfall of memory, and the value is greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that memory is a constraint on this system. Finally, if Column 7 shows a shortfall of memory, but does not exceed the aforementioned limits, the implication is that the system has reached memory saturation, but is still able to process the current workload.

Step 2 - It should be stressed that the value reported in column 7 is calculated over the entire measurement period and therefore could be biased by periods of heavy or light activity. It is for this reason that the user is urged to run the monitor during those periods of time where the workload is considered to be heavy and constant. In order to determine if the above type of biasing is occurring, the user may want to check Plots 1-3. If it appears that there is a mixing of light processing and heavy processing the user may want to re-run the data reduction program, using the time-frame option, to separate the heavy processing time.

Step 3 - Calculate the ratio of column 5 divided by column 4. This is an indication of the maximum number of user jobs that your system can support at any one time, without the occurrence of significant swapping. If the value in column 10 is equal to or exceeds this ratio, then the implication is that the system has reached memory saturation. If the value in Column 10 is within 2 units of the ratio, then the system probably has sufficient memory but is approaching saturation. Finally, if the value in column 10 is less than the ratio by more than 2 units, the current system has sufficient memory.

This step can be further verified by checking columns 14, 16 and 18 for indication of significant swapping.

Step 4 - If column 13 is less than 85, the current system should have sufficient memory and the other steps should not be showing indications of memory problems. If column 13 is between 85 and 95 then the current system is approaching saturation and at times may be showing some indications of a backlog. If the figure exceeds 95, then other steps should be indicating signs of moderate to severe memory problems.

It is possible for column 13 to indicate that sufficient memory is available, while other reports tend to indicate a memory contention problem. This anomoly results from the presence of zero urgency jobs. The *MUM Data Reduction Program* discounts zero urgency jobs that are in memory. Such jobs are treated as if they did not exist. The reason for this is that GCOS is supposed to swap such jobs out of memory if their memory can be used by waiting jobs (see memory algorithm description in subsection 6.3.15). In reality, however, this does not always appear to happen. Therefore, the following condition may arise during data reduction. Assume that a site has 100K of memory configured and that that memory is currently being occupied by 2 jobs -- each of 50K. However, one of those two jobs has a zero urgency. The Data Reduction Program will claim that memory is only 50 percent utilized. However, at the same time, it is possible that several jobs may be waiting for memory, because GCOS is unable or unwilling to swap the zero urgency job from memory. The existence of this problem is highlighted with the Zero Urgency Job Report (subsection 6.3.16). In addition, there is a data reduction input option (ZERO) which tells the Data Reduction Program to include zero urgency jobs in all calculations.

Step 5 - If column 8 is greater than or equal to 3, the indication is that memory has become a constraining factor.

Step 6 - If Column 12 is greater than 2, the indication is that memory wait time is high and that memory is probably a constraining factor.

At this point, the user should have a fairly good indication as to whether or not memory is a constraining factor. The following steps will indicate some additional reports that the user should reference to determine those jobs that might be causing the memory problem.

Step 7 - One of the largest users of resources are jobs that abort and then must be rerun. Aborts usually occur due to user errors, but hardware aborts are not uncommon. If management is aware of aborting jobs and the reasons for them, they can possibly save substantial system resources. The Abort Report is described in subsection 6.3.8 and gives an indication of the system resources being wasted by aborting jobs.

Step 8 - It is important for management to be aware of jobs that are either misusing system resources or are requesting large amounts of system resources. Upon identifying such jobs, these jobs could be redesigned, scheduled for non-peak processing, or, in the case of wasted resources, the waste could be eliminated. The Excessive Resource Report allows the user to uncover jobs of this type and is described in subsection 6.3.10. When using this report, the following are suggested parameter values:

wasted core - 10K
memory - either 35K (WWMCCS standard) or 2 times the value in column 4
CPU time - 15 minutes
IO time - 30 minutes
URG - 40
RATIO - 2

Step 9 - By examining the System Program Usage of Memory Report, the user can determine those system type jobs that are requiring memory. It is possible that some of the system jobs can be eliminated or at least reduced in size. This is especially true for the TSS. However, it must be realized that a limitation on Time Sharing size may adversely effect TSS response. In many cases, if large file transfers are being processed during prime time, the size of the FTS WIN subsystem can rise to 70 or 80K. By not allowing WIN file transfers to run during prime time, significant memory savings can result.

Step 10 - As is explained in great detail in subsection 6.3.15, it is vitally important that the overall urgency level of jobs being processed remain low. The Distribution of Urgency Report can be used to determine the overall urgency level of jobs being processed. This report should show that 60 percent of the jobs being processed at any one time have an urgency level below 40 and that a substantial proportion of these should have an urgency level between 5-10. The summary at the bottom should indicate that 75-80 percent of all activities processed had an urgency level below 20.

If this report indicates a large percentage of high-urgency jobs, then the SNUMB/IDENT report, or the Excessive Resource Report, can be used to identify those particular activities processing with a high urgency.

Step 11 - If the analyst wants to track the memory performance of a given set of jobs, the use of the SPECL input option and the generation of the Special Job Memory Reports will provide sufficient data for detailed memory tracking. This procedure is especially useful in analyzing the memory requirements of TS1, FTS and the special JDA-developed software (JDSIP, JDSUP). Refer to subsections 6.1.24 and 6.3.14 for complete descriptions of these Special Job Memory Reports.

Step 12 - Another indication of poor system performance possibly caused by memory shortfall, tape drive shortfall, poor operator performance or a poor system scheduler design is the long delay of jobs as they pass through the various allocation phases prior to core allocation. The Allocation Status Report, the System Scheduler Delay Time Histogram and the Delay Time Until Core Allocation Histogram can all be used to determine which jobs, and how many jobs, are being significantly delayed during the various allocation phases. These reports are all fully discussed in section 6.

Step 13 - The four time plots should be examined in order to determine if memory problems occur during specified times of day. If this appears to be the case, then an adjustment of the scheduler classes, or manual control of the scheduling of jobs, may alleviate the problem.

Step 14 - In order to perform a valid operation, measurements should be made over several days. Figure 14-6 is a summary check sheet that can be used for this evaluation.

Step 15 - Memory problems may also be occurring as a result of jobs being delayed due to CPU constraints or I/O constraints. In these cases, jobs tend to sit in memory due to a lack of other system resources. Because these jobs are being delayed, other jobs cannot enter memory, and memory demands begin to backlog. Therefore, if memory is a constraint, the user should consider conducting a CPU analysis as well as an I/O analysis.

14.6.4 CPU Evaluation. The CPU evaluation will determine the general utilization level of the processor and then determine if the CPU is dominated by GCOS or user program execution. In addition, the CPU evaluation can be used to determine if jobs are being significantly delayed by a lack of processor power. A CPU data reduction is required for this evaluation. It is also beneficial to have an associated MUM data reduction available for the same time period.

14.6.4.1 Data Recording. The heading page of the CPU data reduction report provides the dispatcher options currently in effect on the system. Recent tests have shown that the Urgency Thruput Option should be enabled, as well as the In-Core Push Area and Dynamic Buffering of SSA Modules. The TSS and the various WIN subsystems should not be placed in Priority B processing. In addition, sites should try to avoid enabling the I/O Thruput Option, unless strong justification exists to decide otherwise. The CPU Time Report is produced every 10 minutes of elapsed time and the data of interest should be found in the last 10-minute report.

14.6.4.2 Evaluating the Data.

Step 1 - The CPU may by considered a bottleneck if the % Idle CPU column of the CPU Time Report is less than 20 or the summary column at the end of the CPU Idle plot indicates that 50% of the time or more, the processor is less than 25% idle. It is a fairly agreed upon standard that average processor utilization should not exceed 80%. By maintaining average processor utilization at this level, the processor will have sufficient remaining capacity to be able to handle those peaks of processor demand that normally occur during the day. The three category breakdowns at the end of the CPU plot represent the three conditions of insufficient power, sufficient power, and excess power respectively.

Step 2 - The % Gate Loop column of the CPU Time Report provides an indication of the percentage of CPU power being lost because multiple

MEMORY STATISTICS REPORT

	#1 DATE	#2 START TIME	#3 STOP TIME	#4 HOURS	#5 CPU/IO RATIO	#6 AVG ACT SIZE	#7 USER MEM AVAL
1							
2							
3							
4							
5							
AVG							

	#8 SYS MEM USED	#9 MEM (+/-) CALC-PALC	#10 AVG USER ACT WAITING MEM	#11 AVG SYS ACT WAITING MEM	#12 AVG USER ACT IN MEM
1					
2					
3					
4					
5					
AVG					

	#13 AVG SYS ACT IN MEM	#14 RATIO DUR/MEM	#15 % SLAVE MEMORY	#16 USER SWAP TIME	#17 ACT WAIT FOR ORIG MEM ALOC
1					
2					
3					
4					
5					
AVG					

	#18 DURATION OF USER ACT	#19 # USER SWAPS	#20 ACT/HOUR (THRUPUT)	#21 USER/ACT/HR (THRUPUT)	#22 SWAP/HR
1					
2					
3					
4					
5					
AVG					

Figure 14-6. Memory Evaluation Check Sheet (Part 1 of 2)

MEMORY EVALUATION

____1. Fill out memory statistics - report form

____2. Review urgency distribution report (most user jobs should be in the 0-20 range)

3. Evaluation of memory statistics:

____a. Calculate 15% of user memory available
#7*.15 = ______ (1)

Calculate two times average activity size
#6*2 = ______ (2)

List the memory surplus or shortfall from #9
Surplus ______ (3)
Shortfall ______ (4)

If (3) > (1) or (2) ⇒ OK
< (1) or (2) ⇒ approaching saturation
(4) > (1) or (2) ⇒ constraint

____b. *Refer to GMF 14-17, step 2, to look for bias conditions, which may affect the results obtained in (a)*

____c. Calculate the ratio of user memory available to average activity size
#7/#6 = ______ (5)

List the average number of user activities in memory
#12 = ______ (6)

If (6) + 2 < (5) ⇒ OK
(6) < (5) ⇒ approaching saturation
(6) > (5) ⇒ saturation

Verify by looking at swap values in
#16, #19, #22

____d. Additional checks for memory problem:

#15 > 85
#10 > 3
#14 > 2

____4. Examine PALC Report, Excessive Report for consistent offenders

Figure 14-6. (Part 2 of 2)

processors are interfering with each other. Within GCOS, there are many tables that must be updated and/or referenced by a processor during its execution. When these tables are being used, the processor must insure that they are not altered in any manner. In order to accomplish this, the processor will lock a gate. This locked gate will prevent any other processor from accessing this table. When the first processor has completed its use of the table, the gate will be opened. If a processor must access a gated table, it will simply perform a "CPU loop" at that table, waiting for it to be opened. The amount of time being spent in this gate locked-CPU loop code is depicted by this column. If this value is greater than 5%, then this is an indication that the multiple processor configuration is beginning to lose its cost effectiveness.

Step 3 - The MUM Excessive Resource Report can be used to determine those jobs requiring excessive CPU resources. In addition, the CPU Plot report can be used to determine those times of day when processor availability is in the critical range. Using these two items of information, system schedular classes can be created based on CPU requirements.

Step 4 - Another indication that the CPU is a bottleneck can be determined from the CPU Queue Length histogram. If the average queue length is greater than two times the number of processors configured, then the processors are being requested to handle an excessive amount of work. This queuing of jobs at the CPU is an indication that during some period of the day, there is insufficient CPU power to handle the workload. Once again, the CPU Plot can be used to determine those periods of time.

In addition, at the bottom of each 10-minute section of the CPU Time Report, a line is printed indicating the CPU queue length during the last 10-minute period, as well as since the start of the run. These figures provide an excellent indication of those times during the day that the processor is being overloaded. Control of the scheduler queues is one method of limiting the amount of work being entered into the system.

Step 5 - The CPU Time Report and WIN Report can be used to determine those periods of time when TSS and/or WIN programs are using excessive amounts of processor time or, on the other hand, do not appear to be requesting sufficient CPU service.

Step 6 - If the Percent of Memory Time in Queue histogram shows that the average activity is spending more than 30% of its memory time waiting for the processor, this is a strong indication that processing power is a constraining factor. Once again, in order to relieve this constraint, it will be necessary to acquire an additional processor, a faster processor, or else the workload will need to be controlled via the scheduler classes. The CPU Access By SNUMB Report can be used to determine, on an activity-by-activity basis, exactly which activities are being delayed the most due to the lack of processor power.

Step 7 - If the CPU Time Report shows that the percent of system CPU exceeds 30%, this is another indication that the system software is being requested to handle an excessive workload. In all probability, the system queues are increasing to such a size that the system software is expending excessive resources managing its queues.

Step 8 - The CPU Time Report (dispatcher queuing portion) indicates the percentage of time the various system programs spend waiting for the processor and the average queue position of these programs. System programs should not spend more than 2% of the time waiting for service and their average queue position should not exceed 2. If this is not the case, the analyst should ensure that the Urgency Thruput Option is enabled. In addition, the Urgency Report and Excessive Resource Report of the Memory Monitor should be examined to ensure that there is not an excessive number of user activities processing at very high urgencies (see section 14.6.3 for memory evaluation details).

Step 9 - Figure 14-7 is a CPU Evaluation check sheet that can be used for this evaluation.

14.6.5 I/O Evaluation. The I/O evaluation will determine whether the mass storage subsystem, or tape channel subsystem is the cause of system degradation. This evaluation requires the user to have processed the Mass Storage Monitor and Channel Monitor data reduction programs. See section 8.6.13 for a discussion of how to limit the processing of the CM data.

14.6.5.1 Data Recording. All output from the Mass Store Monitor and Channel Monitor are required. No individual work tables are required, but the user may generate some if he feels that it will help in his analysis (see figure 14-8).

14.6.5.2 Evaluating the Data. Chapters 7 and 8 provide a fairly detailed description of the procedure to be followed in analyzing the associated reports. In this section, reference will be made to those chapters indicating actual data values that should be used as a reference for comparison.

Step 1 - Read subsection 7.3 and subsections 8.2 and 8.3.

Step 2 - Check the crossbar configuration using the procedure described in subsection 8.2 (see figure 14-9 for sample check sheet).

Step 3 - Examine the Proportionate Device Utilization Report produced by either the MSM or CM. Check for devices which have significantly higher utilization than other devices in the system. These devices are potential bottlenecks and should be more closely analyzed. It is desirable, even though perhaps not possible, to have equal utilization across all disk packs. Read subsections 7.5.18 and 7.5.20 for further details on this step. Once a pack(s) is identified, further analysis should be performed to determine the actual files being referenced on the pack (see step 9).

CPU EVALUTION

____1. Fill out CPU statistics report form.

____2. Review CPU time report. The percent of time in queue and average queue position of system jobs should be less than 2.

3. Evaluation of CPU statistics:

____a. If #5 20 or #6 50 = CPU bottleneck

____b. If #9 5% = multiprocessor configuration is losing cost effectiveness

____c. Calculate number of processors *2 = ______ (1)

If #10 (1) = insufficient CPU power to handle workload

____d. If #11 or #12 30 = insufficient CPU power to handle workload

4. Dispatcher Options:

____a. Urgency Thruput

____b. I/O Thruput

____c. MCOUNT On

____d. Priority B On

____e. In-Core Push Area

____f. Dynamic Buffering of SSA Modules

____5. Time of day of heaviest CPU usage.

____6. Five-ten jobs most affected by CPU queue wait -- look for trends (IDENT/USERID from MUM)

Figure 14-7. CPU Evaluation Check Sheet (Part 1 of 2)

CPU STATISTICS REPORT

	#1 DATE	#2 START TIME	#3 STOP TIME	#4 HOURS	#5 % IDLE CPU	#6 % 25%
1						
2						
3						
4						
5						
AVG						

	#7 % 50	#8 % 50%	#9 % GATE LOOP	#10 AVG QUE LENGTH	#11 % MEM TIME IN QUE	#12 % SYSTEM CPU
1						
2						
3						
4						
5						
AVG						

Figure 14-7. (Part 2 of 2)

I/O EVALUATION

____1. Run FMS SPUTIL Report and FRAG Report. Refer to GMF Manual 14-23, and fill in the disk space utilization form.

____2. Review the Channel Monitor Histograms for I/O queue lengths and I/O queue times. Note any devices which had average queue lengths 1 or average queue times 15 ms. These should be investigated for device contention and proportionate device utilization.

____3. Review the Channel Monitor System Summary Report.

  ____a. Determine the distribution of connects across device types:

  ____b. Determine the percent of connects issued at the third level for each multiple device channel:

IOM #	CHAN #	3rd LEVEL %				
		1	2	3	4	5

____4. Fill in I/O Statistics Report form from MSM and CM reports.

5. Evaluation:

  ____a. If #5 55%, should try to increase FMS Cache.

  ____b. If #9 25K, then large subsystem user may be degrading TSS response. Review CAM to verify.

  ____c. Review #6, #7, #8 to determine if AST Buffers are adequate:

  ____d. If #10 90, SSA Cache needs to be increased.

  ____e. If #11 7%, examine system modules which need to be moved to SSA Cache.

____6. Review the Mass Store Monitor Seek Movement Histograms for potential problem devices

____7. Review the Mass Store Monitor Head Movement Efficiency Report. Efficiency of 1.5 indicates potential problem.

____8. Explain how to get SSA Cache hit ratio and FMS Cache hit ratio using PEEK command from console.

Figure 14-8. I/O Evaluation Check Sheet (Part 1 of 2)

I/O STATISTICS REPORT

	#1 DATE	#2 START TIME	#3 STOP TIME	#4 HOURS	#5 FMS CACHE HIT RATIO
1					
2					
3					
4					
5					

	#6 NIAST EFF	#7 NIAST BUFF SUFF	#8 NIAST DELAY RATIO	#9 TSS SWAP TRANSFER SIZE (80%)	#10 SSA MODULE BUFFER HITS
1					
2					
3					
4					
5					

	#11 SYSTEM I/O%	#12 % WRITE VERIFY	#13 % CONNECTS ACROSS 181	191	451	500	501
1							
2							
3							
4							
5							

	#14 5 MOST HEAVILY USED DEVICE IDS	#15 HEAD MOVEMENT EFF 181	191	451	500	501
1						
2						
3						
4						
5						

Figure 14-8. (Part 2 of 2)

MPC/PSIA CROSSBAR CHECK

____1. List the IOM channels on the crossbar card

____2. From the MPC card, list MPC/PSIA for each channel

3. Problem areas:

   ____a. Same MPC/PSIA appears consecutively (logical channel problem)

   ____b. Same MPC number and PSIA 0,1 and 2,3 appear consecutively (link adapter conflict)

IOM/CHANNEL	MPC/PSIA

Figure 14-9. Crossbar Check Sheet

Step 4 - The histogram displaying Data Transfer Sizes for TSS Swap Files can give a strong indication of the sizes of TSS subsystems being used. TSS subsystems of over 25K can cause significant increase in overall TSS response, especially if several of these subsystems are being executed simultaneously. If more than 20 percent of the entries in this report fall in the bucket ranges above 25000, this is a strong indication that TSS response might be a problem. This problem can be further confirmed with the CAM.

Step 5 - Seek Elongation - Subsections 7.5.6, 7.5.7 and 7.5.8 describe in detail the procedures used to investigate seek elongation problems. An average seek of over 50 cylinders for DSS191s and 100 cylinders for DSU450s and DSU500s should be considered significant. If seek elongation problems are discovered, further analysis should be performed to determine the actual files being referenced on the pack (see step 9).

Step 6 - Analyze the Channel Monitor Idle Report. This report can be generated only if the Idle Monitor was run in conjunction with the Channel Monitor. If the "% of Idle Time During Which I/O Was Active" value exceeds 25%, then substantial benefit may be obtained by eliminating I/O contention. The above value is an indication that even though the CPU is going idle (i.e., has no useful work to perform) there really is potential CPU work available. However, under current conditions, this potential CPU work is being delayed because of I/O contention.

Even though the above figure exceeds 25%, the system may not have sufficient CPU power available to handle the increased work generated by removing the I/O contention. Therefore, the analyst should also check that the "Average System % Idle" figure exceeds 15%. If this proves to be the case, then removal of any I/O contention should prove beneficial. On the other hand, if the figure is lower than 15%, then removal of any I/O contention will probably result in additional CPU contention. The Idle Report will also indicate those devices causing most of the contention. Make a record of the device numbers.

Step 7 - Examine Channel Monitor reports for I/O queue time problems. In performing this step, the following reports should be used:

a. Channel Busy and Device Busy Report. This report would indicate that both a channel shortage and a device contention problem exist. In order to alleviate the problem, files would not only need to be removed from a given device, but also moved to a new device located under a different channel subsystem. If additional channel power was acquired, then files could be removed from the device exhibiting contention and moved to another device under the same channel subsystem. An entry should be considered a candidate for further analysis if more than 20 percent of all its accesses are queued.

b. Channel Busy and Device Free Report. This report would indicate that accesses to a given pack are being delayed because of a lack of channel power. The solution to this problem is to increase the channel capacity of the system, or else move a significant number of files to devices located under a different channel subsystem. Any device which has over 20 percent of its accesses being delayed because of channel contention should be considered a candidate for further analysis.

c. Channel Free and Device Busy Report. This report would indicate that a given device is a potential bottleneck if over 20 percent of its accesses are being queued. The solution to this problem is a relocation of files off of the device in question.

d. Device Free But Has a Queue Report. This report shows the number of times a connect was made to a device, the device was free (no active connect was in process), but yet there were outstanding requests waiting for the device. This report is another indication of a channel contention problem (similar to Channel Busy and Device Free Report). The analyst should refer to subsection 8.5.6.9 to see the differences between the two reports. An entry should be considered significant if over 20 percent of its accesses are delayed in such a manner.

e. 500 Disk Drive Report. The analyst should refer to subsections 7.5.7 and 8.5.6.10 for an explanation of the physical characteristics of a 500 device and the information to be obtained from this report.

Step 8 - If certain devices have been determined as bottlenecks under the procedures described in step 7, the Job Conflict Report should be obtained for those devices following the procedures described in subsection 8.5.12.

Step 9 - Execute the Mass Store Monitor Data Reduction Program. Following the procedures described below, the analyst should be able to determine the exact files that are causing the contention found under the earlier steps:

a. Review the System File Use Summary Report and the Individual Module Activity Report (subsections 7.5.9 and 7.5.10) to determine whether system files, SYSOUT files, accounting files, or TSS swap files are located on the device in question. If accesses to these files are a significant percentage (20 percent) of all accesses to the device, then relocation of these system files may alleviate the problem.

b. Using the procedures described in subsections 7.5.12, 7.5.13, 7.5.16 and 7.6.1, determine all files that were accessed on the devices in question.

Step 10 - Using the CM, it is possible to perform a detailed analysis on channel queuing for a particular job. Details for this procedure can be found in subsections 8.5.13, 8.5.14 and 8.6.10.

Step 11 - This step outlines procedures for relocating files identified as candidates for file relocation. Because of automatic load-leveling activity by the GCOS operating system, an analyst has only limited flexibility for the placement of system, permanent, and temporary files:

a. System Files. The device name on which a system file is to be placed can be specified at system startup. Care should be taken to insure that multiple high-used system files are not placed on the same disk device. If possible, separate high-use system files across disk subsystems. In addition, ensure that SSA Cache memory and FMS cache are enabled to reduce disk I/O activity to certain system files. Details for this analysis can be found in subsections 7.5.9, 7.5.10 and 7.5.19.

b. Permanent Files. The device name for a permanent file can be specified at creation, whether through FMS or the ACCESS subsystem of Time Sharing. Files can be moved by changing their names, creating a new file with the old name, and moving the data. The new file can be created with a DEVICE specification.

c. Temporary Files. The device name for a temporary file can be specified in the second field of the $FILE card in the job control deck. Jobs which run frequently can have their $FILE cards changed. Other jobs can be controlled by policies governing the use of $FILE cards.

Additionally, sites that have different device types may specify preferred device types to be used for temporary files. This procedure will allow activities requiring disk storage to take advantage of higher speed devices.

Step 12 - This step identifies possible seek contention problems attributable to inadequate temporary file space. This procedure uses the Disk Fragmentation Report (FRAG) available at most sites. If such a report is not available, contact CCTC/C751. It is necessary to analyze temporary disk capacity on all disk units rather than just the units identified in previous tuning steps. This analysis is necessary because the disk units exhibiting high activity due to temporary file use often have more available temporary space. The increased utilization of these disk units may be caused by inadequate temporary storage on other disk devices. For this analysis a form as shown in figure 14-10 may prove useful.

a. Report Values. The FRAG Device Report contains the following information on each disk device: (1) device identification, (2) overall capacity, (3) available disk space, (4) disk space dedicated to permanent storage, (5) number of disk fragments, (6) average fragment

DEVICE	TOTAL CAPACITY	TEMP CAPACITY	% AVAIL	RATIO	NUMBER OF FRAGMENTS	MAX SIZE	PERCENT LLINKS 1-12	PERCENT LLINKS 13-120	TOTAL PERCENT	FRAGMENT

Figure 14-10. Temporary Storage Test Form

size, (7) maximum fragment size, (8) percentage distribution of fragments by size, and (9) total fragmented space.

b. Form Entry. Enter the device identification for each disk device on the Temporary Storage Test Form. For each disk device enter: (1) the LLINK capacity column; (2) the temporary available; (3) the available percent; (4) the number of fragments in the Number of Fragments column; (5) the maximum fragment size in the Max Size column; (6) the percentage of llinks (1-12 and 13-120, respectively) in the Percent Fragments column.

c. Calculations. Add the Percent Llinks 1-12 and the Percent Llinks 120 columns and place the sum in the Total Percent column.

d. Decision. Place a check mark in the "Ratio" column if the available percent is less than 16 percent. Place a check mark in the "Fragment" column if (1) the number of fragments is 100, (2) the maximum size of a fragment is 2000, and (3) the total percentage of fragment llinks less than 121 LLINKS (10 LINKS or less) is 40.

Step 13 - If devices have been checked in Step 6, then the temporary file space on these devices is constrained by either (1) insufficient file space or (2) disk fragmentation.

If a large number of devices need to be checked, then overall temporary disk space availability may be a constraint to system performance. At this point a site should either institute procedures to recover permanent disk space (purging of unused files) or re-evaluate disk capacity relative to system workload.

If certain disk devices have been checked because of their temporary to total ratio, or if many devices have been checked because of disk fragmentation, then the following procedures to balance the file system should be instituted.

a. Fragmentation. Temporary disk space may be compacted by a full restore of the file system (cold boot). The Disk Fragmentation Report, run monthly, could help determine the frequency at which a full file system restore would be necessary.

b. Amount of Available Space. The full file system restore redistributes permanent files to equalize the amount of temporary storage on each pack. The FMS SPUTIL Report should be examined after the restore to verify that disk units have been balanced. Temporary disk space may still not be evenly distributed if (1) users have specified specific devices for permanent files or (2) unusually large files are present on certain disk units. Consider such factors as intentional file placement for maximum seek activity and operational requirements of data bases before moving permanent files.

14.6.5.3 Mass Storage Operation. The following article appeared in the September 1981 WWMCCS H6000 World Magazine. It presents an excellent description of disk operation and some problem areas that a system analyst should be aware of.

"We won't discuss the gyrations that GCOS performs in setting up an I/O request; e.g., management of PAT pointers, PAT bodies, etc. (We recommend the HIS GCOS Analysis course). Our interest begins at the point when an I/O request generated by some workload element is presented to the I/O Supervisor (IOS) for service. We don't care whether the I/O came via GFRC or directly from a MME GEINOS. We do care about the major logical/physical events from this point to I/O termination which are important to understanding system performance.

"As an example, let's assume that we have a dual-channel micro-programmable controller (MPC) cross-barred to its disk devices. Dual-channel means that the MPC has two physical data paths for transferring data to the IOM. Cross-barred to the device means that there are two physical paths from the MPC to each of its devices.

"Let's also assume that the IOM is cross-barred at two-to-one. IOM cross-barring refers to the logical-physical hardware channel relationship of the IOM. This relationship is defined to GCOS in the STARTUP deck and maintained in the Secondary Configuration Table and elsewhere. For the example, assume that PUB addresses 0-8 and 0-9 are associated with one physical MPC channel, and that 0-10 and 0-11 are associated with the second MPC channel as is shown below:

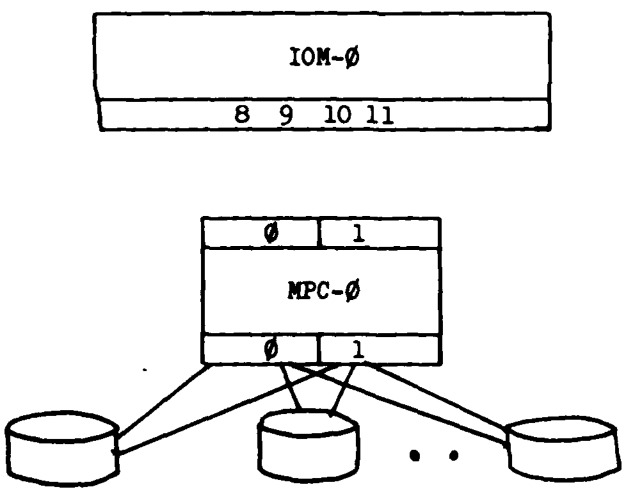

Figure 1

"Defined in the STARTUP deck is the primary PUB address for the disk subsystem and the order in which logical paths will be assigned by GCOS during processing. The card

```
$       IOM-0           PUB-8,MS0450,UNITS-6,
$       ETC             UNIT-1,ST1
          :
```

defines PUB 0-8 as the primary PUB address for a MSS451 subsystem. The card

```
$       XBAR            IOM-0,PUB-8,PUB-10,PUB-9,PUB-11
```

defines the order in which PUB assignments will be made. This information is maintained by GCOS in the I/O stream table. All requests for a device are made with its primary PUB address, e.g., an I/O request for device 6 on this subsystem would request IOM-PUB-DEVICE = 0-8-6.

"Suppose that IOS receives a request for I/O to 0-8-6. If the device is currently performing an I/O (i.e., busy), the I/O request is queued for later service. If the device is free, IOS will assign a logical channel to the request if one is available. It will attempt to assign PUB 0-8. If 0-8 is busy, it will attempt to assign 0-10, etc., in the order prescribed on the $ XBAR card. If no logical channel is available, the I/O request will be queued for later service as before.

"To recap, for an I/O request to become active (in-service), both the requested device and a logical channel associated with the MPC which controls the device must be free. Otherwise the request is placed in the I/O queue. Normally the request will be placed at the end of the queue; i.e., behind earlier I/O requests. However, certain I/O requests are linked in front of the queue, e.g., TSS requests for access to its swap files (#S, #T, #U, #V). This can impact system performance and will be discussed later.

"Whether or not our request for I/O to 0-8-6 becomes active, IOS will scan the I/O queue to determine if any earlier I/O requests can now be activated (device and logical channel free). When an I/O request can be activated, IOS loads the mailbox in Hard Core Memory for the logical channel and signals the IOM. This mailbox contains information required by the IOM to process the I/O independently of the mainframe processor. The IOM loads the mailbox information into its scratch pad memory and begins processing the physical I/O.

"Thus far, the events which have occurred were independent of the type of disk device involved. The three phases of physical delay for a disk device are seek time, rotational latency, and data transfer. The IOM issues a seek request to the MPC for 0-8-6, but now performance implications arise which depend both on the physical characteristics of the device, and on device-dependent interactions with IOS which might be required.

"First let's look at seek time, defined as the time required to move the read/write heads from their current position to the cylinder requested by the current I/O. For a given device, seek time is proportional to the number of cylinders traversed. Table 1 gives seek times for H6000 disk devices. The table provides sufficient information to derive the average seek time for a device from the average number of cylinders moved (available from MSM).

Device	DSS181	DSS191	MSU450
Minimum	10 ms	10 ms	8 ms
Average	34 ms	30 ms	25 ms
Maximum	60 ms	55 ms	55 ms
Cylinders/pack	200	404	808

Seek Specifications
Table 1

"Now that 'seek' has been defined, we can discuss the motivation of implementing the logical-physical channel architecture. The disk subsystem can have as many seeks and data transfers ongoing as there are logical channels assigned to the subsystem. The number of simultaneous data transfers that can be performed is limited by the number of physical channels. In our example, we could see two devices connected to the physical channels transferring data and two additional devices moving the read/write heads to the requested cylinder (seeking). Or, we could have one data transfer ongoing with three seeks, or we could have four seeks ongoing, etc. So the purpose of IOM cross-barring is to increase the simultaneity or overlapping of device seeks.

"When the seek to the requested cylinder has been completed, the device is ready to transmit or receive data. Here again a device-dependent performance consideration comes into play. For a DSS191 or MSU450 disk subsystem, once IOS has loaded the channel mailboxes and signalled the IOM, it performs no further service to the I/O request until an I/O termination interrupt is received for that device. The determination of seek completion and assignment of the physical channel for the data transfer takes place completely in the MPC. Not so for the DSS181 subsystem.

"The DSS181 does not provide a seek complete interrupt, and so IOS must check for a seek complete status each time it gets control for an I/O request as described above or for servicing an I/O termination interrupt. This introduces another delay for the 181s from the time seek completes until the next time IOS is activated. When IOS detects seek complete for a 181, it must issue another command to connect the physical channel and begin data transfer.

"If both of our physical channels are busy transferring data, there is another delay until a physical channel is free. But let's suppose the seek has completed, a physical channel is available for connecting to the device (whether by the MPC or by IOS) and we are ready to transfer data. Another delay is imposed now due to the rotational latency of the device. This is the time from physical channel assignment until the requested sector rotates to the read/write heads. This delay is a function of the rotational velocity of the device. The minimum latency for a device is always 0. (The data is at the heads at channel connect time.) Maximum latency for a device is the time required for 1 revolution of the disk, and average latency is generally defined as the time required for 1/2 revolution (see Table 2).

Device	DSU181	DSU191	MSU451
Rotational Speed	2400 RPM 40 RPS	3600 RPM 60 RPS	3600 RPM 60 RPS
Avg. Latency	12.5 ms	8.3 ms	8.3 ms
Max. Latency	25 ms	16.7 ms	16.7 ms

Rotational Latency
Table 2

"A word about Rotational Positional Sensing (RPS). RPS is a feature available on some disk subsystems -- a required option (?) on MSU451s -- which keeps the MPC informed about the rotational position of each disk. In a situation where the physical channel is free and more than one device has completed it's seek, the MPC can assign the physical channel to the unit whose requested sector is closest to the heads. The idea is to reduce rotational latency thereby increasing the physical channel time which can be utilized for data transfer. What does RPS buy you if only one disk unit has completed its seek when the physical channel is assigned? Not a thing.

"Now that device 0-8-6 has been assigned the physical channel and the requested sector is under the heads, the physical transfer of data begins. The time to complete the transfer of data depends on the number of words to be transferred, the rotational velocity of the device, and the recording density of the media. Data transfer rates per second are given in Table 3.

Device	DSU181	DSU191	MSU451
Transfer	416KC	1074KC	1074KC
Rate	69.33KW	[illegible]9KW	179KW

KC = 1000 - 6 BIT CHARACTERS
KW = 1000 - 35 BIT WORDS

Data Transfer Rates
Table 3

"When data transfer has completed, an I/O terminate interrupt is issued to the H6000 CPU (CPU 0). The GCOS dispatcher suspends execution of the program in operation, performs a few housekeeping duties and then dispatches CPU 0 to the disk channel module in IOS. A bit more housekeeping and, for our purposes, the I/O to 0-8-6 is complete. IOS then checks the I/O queue for seeks to start -- note that at this point it has a free logical channel, and so can start at least one seek if an I/O request is pending against a free device on this subsystem. If a DSS181 subsystem is present, it must also check for a seek complete status as described above. When IOS can perform no more actions, it relinquishes control to the dispatcher.

"In terms of capacity, a DSU1191 can hold as much data as approximately 4.3 DSU181 packs (see Table 4). A MSU451 can store roughly 2 DSU191 packs or 8.6 DSU181 packs worth of data. In terms of data transfer rates, the 191s and 451s are about 2.6 times faster than 181s.

Device	DSU181	DSU191	MSU451
Char. per pack	27.6mc	117.9mc	235.8mc

Device Capacities
Table 4

Here are a few WATCHITs.

"WATCHIT 1: When replacing 181s with 451s. The vast capacity of the 451s make it feasible to replace up to 8 181s with a single 451. In a situation where 4 I/Os were pending to files residing on 4 different 181s (our 2 to 1 cross-barred dual channel example), as opposed to the 451, the 181 configuration might be preferable. Although the 451 is 2.6 times as fast, we would sacrifice overlapping of the seeks, and the 181 subsystem can transfer data from 2 of the devices in parallel. The CM Device I/O Queue Length Histrogram and the Channel Free-Device Busy Report can indicate if a device contention problem exists.

"WATCHIT 2: When pondering vendor specs for 181s. We modified the MSM data reduction program to obtain the average data transfer size to each disk device. MSM also provides the average cylinders seeked per I/O for each device. Assuming the average rotational latency, one should be able to use the vendor specs provided in this article to compute the I/O service time for a given device. The I/O service time for each device is provided as a CM report. For MSU451s the computed value matches the CM report to within a few milliseconds. However the 181 computed and monitored values differ by as much as 30ms and more -- the computed value always being less. Since channel utilization is fairly low, the most likely reason for this discrepancy is that IOS latency time we discussed earlier.

"WATCHIT 3: When placing TSS swap files (#S, #T, #U, #V). We mentioned earlier that TSS requests for access to it's swap files were linked to the front of the I/O queue. TSS is not one to wait on user I/Os, nor does it break it's swap action into nice little 320-word standard system format I/Os. It's gonna swap the whole thing. Say, there's a nice little TSS COBOL sub-
system -- only 40K! Time to swap? So soon? Oh yeah, that memory-time quantum. Lot of memory -- not much time. Oh well, a 40K transfer is OK. My swap files are on fast 451s. Friend, you just tied up that device for almost a quarter of a second, and more importantly, a physical channel as well. That's only half of it. TSS swapped it out, you can bet it'll swap it back in. So we're up to .5 sec of physical channel time to free up some TSS user memory. This scenario impacts total system throughput, not just TSS. That's not all the havoc that larger user subsystems create, but that is enough heartburn for now."

14.6.6 Communication Evaluation. The communication evaluation will determine the overall terminal usage of a system. It can also be used to examine the DN355 usage. This evaluation can be done using either the CAM or the GRTS monitor, or both. Figures 14-11 through 14-14 are sample table formats that may be used to display the gathered data.

14.6.6.1 Data Recording. For figure 14-11, the Terminal Session Report is used. All users with TSS subsystems of 35K are recorded. Column 7 is obtained from scanning the Excess Think/Response Time Report. Figure 14-12 is obtained from the Response Time Report. All periods of time when the response for TSS is greater than 15 seconds are recorded. Figure 14-13 is obtained directly from the High Terminal Usage Report. All DAC terminals with over 90 percent usage, except WIN lines, are recorded.

Figure 14-14 is obtained from the H6000-DN355 Reject Report and the Abort Report. The H6000-DN355 Reject Report is used if the average number of

#1	#2	#3	#4	#5	#6	#7
DATE	TERMINAL ID	USERID	LOGON TIME	LOGOFF TIME	TSS SUBSYSTEM SIZE	NUMBER OF EXCESSIVE RESPONSES

Figure 14-11. Large TSS Subsystem Users

DATE	TIME	RESPONSE TIME	IN-RANGE RESPONSE	NUMBER OF OUT-OF-RANGE RESPONSES	NUMBER OF USERS

Figure 14-12. Poor TSS Response Log

DATE	TERMINAL ID	PERCENT USAGE	NUMBER SESSIONS	NUMBER INPUTS	NUMBER OUTPUTS

Figure 14-13. High Usage Terminals

DATE	TERMINAL ID	OPCODE REJECTS	PERCENT REJECTS	NCP DISCONNECTS

Figure 14-14. Terminal/Line Error Report

reject commands per hour of running time is greater than 50 (total number of reject commands/number of hours in run). Terminals with more than 30 percent of the rejects should also be listed. For the Network Control Program (NCP) disconnects, all NCP 01 line disconnects listed in the Abort Report should be tallied.

14.6.6.2 Evaluating the Data. The following procedure should be followed in order to analyze the data.

Step 1 - TSS response time is dependent upon certain TSS parameters. The TS1 Initial Parameter Report contains current settings of these parameters. The critical parameters are:

a. Initial TS1 Max Size - If operators must increase max size during the day, TSS can slow its processing of user responses until it can grow. This should be set to the normal maximum size TSS reaches.

b. Size Growth/Reduction Factor - These sizes should be identical or growth can be twice reduction.

c. Max Number of Terminals - Should be large enough to satisfy all possible users.

d. Large Subsystem Size/Wait Time - The average subsystem size should be less than 35K and all large users (greater than 35K) should be penalized.

e. Number Swap Files - On an active TSS system, should be four.

f. Allocated Devices - Swap Files (#S, #T, #U, #V) should be on specific devices, as determined by MSM.

Step 2 - Users of large TSS subsystems cause TSS to swap other users out and therefore generate extra TSS and system I/O overhead. If an excessive number of users are using subsystems of greater than 35K (figure 14-11), these users should be queried as to what subsystem they are using. Any large subsystem that has high usage should be investigated for possible rewrite.

Step 3 - The overall system and TSS response time should be monitored using figure 14-12. Periods of bad response should be checked to ensure the bad responses are not caused by a few (less than 10 percent of all in-range responses) out of range responses. If TSS response is truly poor, a correlation between response and large subsystem usage should be attempted (figure 14-11 and figure 14-12).

Step 4 - Terminals logged onto the system for long periods of time (greater than 75-80 percent of the monitoring session), but having few inputs and/or few inputs/outputs, should be investigated (figure 14-13).

Terminals with few inputs, but many outputs, are probably logged onto VIDEO. The number of users on VIDEO should be restricted to one or two per DATANET, due to the buffer load place on the DATANET by VIDEO. If more users are required, monitor output terminals should be used instead of more VIDEO users. Terminals with few inputs and few outputs may be being used just to keep a terminal logged on. This practice causes unnecessary TSS size and processing overhead.

Step 5 - Terminals with a large number of Opcode Rejects (figure 14-14) are an indication of possible line problems. The terminals and lines should be checked for noise and transmission errors. Numerous NCP disconnects per day (figure 14-14) is usually an indication of line or IMP problems. NCP should have fewer than five disconnects per day.

14.6.7 Timesharing Evaluation. The TSS evaluation will attempt to determine causes for poor TSS response time periods. This evaluation can be done using the TSS Monitor of GMF.

14.6.7.1 Background. The GMF/TSSM data collection tool is one of several GMF monitors designed to collect performance information from major portions of the GCOS operating system and related software components. The TSSM is designed to monitor the performance of the TSS executive software. One hundred and four data collection points have been established within the TSS executive code, tracing the execution of nearly every major routine. These tracing points are sufficient to allow for the isolation of time sharing response time problems to, at most, a few routines.

The trace points divide the processing performed by the TSS executive into several states. These states or processing subdivisions allow for the nature of a delay to be identified. The states described are:

a. Eligible for Processor Dispatch. During this state, the user (subsystem) is not being processed by the TSS executive, but is under the control of the system dispatching module, .MDISP.

b. User File I/O. During this state, the user is waiting on the completion of a file I/O operation which the user (or his subsystem) is responsible for initiating.

c. Command Processing. The user cannot affect the time spent during the scanning, analysis, and subsystem loading caused by his entry of time sharing commands, but the overhead incurred due to these actions is directly attributable to him.

d. Derail Processing. In order to receive service from the TSS executive (and from GCOS) a user subsystem uses DRL instructions. When, during execution, a DRL instruction is encountered, a fault occurs, and the user is removed from the subdispatch queue and placed in a fault queue. When this fault queue entry is serviced, the user may enter the "Derail Processing" state to service some subsystem request. As this is the major cause of fault queue entries, this state will track the time required to service the derail request, and to return the user back to the first state defined above.

e. Processor Allocation. During the processing associated with this state, TSS examines the outstanding requests for fault queue service, memory allocation, or line service. This state is ended when the allocation routine locates the correct servicing routine and transfers to it.

f. UST Management. In order to keep track of each of the TSS users, the executive maintains a User Status Table for each terminal which is active. The processing which occurs within this state involves the management of this list of UST entries. Entries may be added, moved, compressed, released, and updated.

g. Memory Management. TSS maintains space for both user subsystems and User Status Tables. Allocation of this space is made based on a complicated algorithm which uses the amount of space required, the amount used, and the priority of various users. TSS will adjust its overall size based on the demands placed upon it, and within the bounds established by both boot deck patches and operator commands. The periods spent by TSS within this logic will be accounted as falling within this state.

h. Swap Space Management. The TSS executive maintains its own swap files for user subsystems. The decision to swap or unswap a user's subsystem is also a complicated process, and involves not only the demand being experienced, but also the time spent by a user without doing I/O. The periods spent by the TSS executive during the management of its swap space will be accounted as falling within this state.

i. PAT Table Management. The TSS system also controls the PAT (Peripheral Allocation Table) space used to define the accesses made by users to temporary and permanent files. While the time spent during this processing is directly attributable to some user action, the importance of this processing indicates a need to discern between the time spent by TSS for other functions and that spent during the management of these PAT Tables.

j. TSS Idle. The time sharing executive may be idle due to an unsatisfied MME request, or due to its having no outstanding work. When either of these conditions arise, the program will relinquish control, and will not receive processor attention until some condition changes.

k. Line Management. TSS employs MMEs to communicate with the GCOS modules DNET and ROUT, which are responsible for communication with the remote terminals used by TSS users. Traces captured while in this executive state will define exactly the response time experienced at each TSS user terminal.

If the boundaries between these states can be traced, and the time spent during a particular interval when poor response is indicated is compared from one state to the next (as well as across common states during different response intervals) the source of a poor response can be isolated, and the probable cause inferred. This is the most that can be gained by a tool of this nature. It should not be expected that the GMF/TSSM will be able to pinpoint the exact cause of all possible response time problems.

The data collected by the existing GMF/TSSM can be useful towards the ends for which it was developed. The data collected by the one hundred and four existing trace points can be used to identify periods of poor response time, and by separating the processing performed by the executive into the states described above, the source of these periods can be isolated and categorized.

14.6.7.2 Report Organization. The time sharing executive is a complex system. TSS contains numerous tunable parameters which can be adjusted by site option patches. When a response problem is encountered, the most likely solution will involve some adjustment of these tunable parameters to better match the workload being experienced at the site. Section 14.6.7.3 of this report presents an overview of the performance considerations relevant to terminal response time, and describes each of the tunable parameters associated with the time sharing executive.

Before response time problems can be analyzed, they must first be identified. While it is probable that no analysis would be performed if poor response were not already being noticed by users, response time anomalies do exist which will not be noticed by even the most observant user. These "invisible" delays can sum to a considerable performance problem. Subsection 14.6.7.4 describes how the user can analyze the reports produced by the TEARS system to identify both obvious and subtle response time and performance problems.

Once a poor response time problem has been identified, it must be analyzed to determine its cause. The isolation and analysis of response time problems based on the reports produced by the TEARS system is discussed in subsection 14.6.7.5.

14.6.7.3 TSS Tunable Parameters. Within the TSS executive, many of the algorithms used to control user processes can be adjusted to better process the particular workload experienced at a site.

The tunable parameters associated with the Time Sharing System have been categorized into several areas, each of which is described in the sections below. The reader is warned that the changes to the default settings for these parameters should be made with care.

14.6.7.3.1 Priority 'B' Processing Parameters. The most important of the tunable parameters associated with TSS response time are those associated with the Priority 'B' processing option located in the system dispatching module, .MDISP. These options will allow the TSS executive to be considered for dispatch with a much higher frequency than the consideration given to normal (batch) slaves. CCTC/C751 does not recommend the activation of these processing options because of the overhead faced by other (batch) jobs when they are active. Activation of Priority 'B' will, for example, limit the processor attention given to WIN system jobs.

There are several parameters which are associated with Priority 'B' processing, and the settings for each should be carefully considered. Many WWMCCS sites have mistaken these parameters for switches which are "turned on" to enable Priority 'B'.

The site option patch shown below will enable Priority 'B' processing for the system.

```
CC      CC   CC                                                       CC
1       8    1                                                        7
             6                                                        3
000000  OCTAL 000001000000 BIT18=1  ENABLE PRIORITY B      (#A1E)     .MDISP
```

Once enabled, the dispatcher will search a list of priority jobs prior to initiating each dispatch. To add TSS to this list of priority jobs, the following site option patch must be used.

```
CC      CC   CC                                                       CC
1       8    1                                                        7
             6                                                        3
000001  OCTAL 6245500000XY  PLACE TSS IN PRIORITY B TABLE (#A1X)      .MDISP
```

'X' is equal to the number of non-class 'B' priority dispatches for each dispatch used by TSS. For example, if this variable is set to one, then TSS could receive every other dispatch. 'Y' is equal to the number of 32 millisecond time units which should be granted on each dispatch to TSS before timer runout faults occur.

Finally, the following site option patch is used in conjunction with #A1X shown above. This patch specifies the processor requests which will be made by TSS.

```
CC      CC    CC                                                CC
1       8     1                                                 7
              6                                                 3
000174  OCTAL MMMMMMNNNNNN  .TAGPP PRIORITY B WORD      (#AOY)  .MTIMS
```

MMMMMM is equal to the number of 32 millisecond time units required for each dispatch to the TSS executive. NNNNNN is used to specify the relationship between TSS and other Priority 'B' jobs. In NNNNNN=1, then TSS will be considered on every other Priority 'B' dispatch. If NNNNNN=3, then TSS will be considered on every fourth Priority 'B' dispatch.

The settings for MMMMMM and NNNNNN should be made in conjunction with particular site requirements. It is recommended that the number of 32ms time units associated with a dispatch be large enough to ensure that TSS is not forced to request a second dispatch to satisfy all outstanding requests. The default value for MMMMMM is 4. The default value for NNNNNN is 1. Thus, without this patch, TSS will request a 4*32 ms time slice, and will be considered on every second priority 'B' dispatch.

14.6.7.3.2 Placement of TSS Files. The TSS executive maintains up to seven files to hold pending deferred sessions, user subsystems, and to swap/push-down active user memory. These files can be allocated each time the time sharing system is started, or they may be declared as permanent. At a minimum, it is recommended that the swap files maintained by TSS be declared as permanent. By doing this, the accesses made to these files by the TSS executive can be optimized by avoiding the problems associated with disk fragmentation.

Most sites have already made their deferred queue file (DQ1) permanent, to allow scheduled deferred sessions to be saved across system boots.

In order to assign the TSS deferred queue and swap files as permanent, the following cards should be placed in the $EDIT section of the startup deck.

```
$       FILDEF      dvc,DQ1,48
$       FILDEF      dvc,SW1,nnn
$       FILDEF      dvc,SW2,nnn
$       FILDEF      dvc,SW3,nnn
$       FILDEF      dvc,SW4,nnn
```

The device assignments for these files (dvc) should be made based on the following criteria:

- Different swap files should not occur on the same device.

- Where possible, the swap files should be adjacent to heavily used GCOS files, thus minimizing I/O times.

Placement of the DQ1 (deferred queue) file is not as critical as the swap files, as it is not as heavily accessed. The sizes of the TSS swap files (nnn) should be made based on configuration limits, but should not be lower than 2260 LLINKS.

In addition to the startup cards shown above, the following site option patches must be included. This patch changes the default USERID for TSS files from OPNSUTIL to GCOS3, thus allowing the permanent deferred queue and swap files to be physically near other GCOS system files.

```
CC      CC    CC                                                  CC
000132  OCTAL 272346620320 DEFAULT UID FOR TSS FILES GCOS3        .MTIMS
000362  OCTAL 202020202020                                        .MTIMS
```

The reader should note that the use of this patch, which changes the default location for TSS files, will also change the permissions requirements for the files associated with the GMF collector. USERID B29IDPXØ must be given modify permissions against GCOS3/TS1, if B29IDPXØ is to be used to install the patches which enable the TSSM. Also, the catalog/file string containing the H* file for the collector program should be modified to give GCOS3 read permissions.

Finally, the device placement for the TSS program files should be made to avoid contention with FMS. Program files should be forced to occur on some device which does not contain an SMC or SMC duplicate. The following site option patch declares the device to use for each of the TSS program files, and associates the name used on the $ FILDEF cards described above with the file codes used by TSS.

```
CC      CC    CC                                                  CC
1       8     1                                                   7
              6                                                   3
000365  OCTAL 400000626601   SWAP FILE SW1 PERMANENT      (#A1D)  .MTIMS
000366  OCTAL 400000626602   SWAP FILE SW2 PERMANENT      (#A1D)  .MTIMS
000367  OCTAL 400000626603   SWAP FILE SW3 PERMANENT      (#A1D)  .MTIMS
000370  OCTAL 400000626604   SWAP FILE SW4 PERMANENT      (#A1D)  .MTIMS
000361  OCTAL 626302         TSS PRIM PGM FIL IN ST2      (#A1D)  .MTIMS
000362  OCTAL 626303         TSS SCND PGM FIL ON ST3      (#A1D)  .MTIMS
```

During TSS initialization, allocation (of the TSS program files) will be attempted on the device specified in the site option patch described above. If allocation is impossible on the specified device, a message will be printed on the system console and allocation will be attempted on another device. When this message begins to appear, site personnel should either schedule a cold boot or attempt to depopulate the device indicated to allow allocation of the TSS file.

14.6.7.3.3 TSS Swap File Processing. If the reader has chosen to assign TSS swap files as permanent, he may disregard the discussions in this section which relate to the growth and release of swap file space. The size (as declared on the $ FILDEF card) for permanent swap files will not change.

The TSS system will maintain up to four separate swap files to hold user subsystems which have been either swapped or pushed-down. When a user subsystem is moved to these swap files, TSS will place it on the least busy of the active swap files. The size of these swap files can be adjusted if enough room is not available on any one of them. The process which will grow and shrink these swap files can be adjusted through the use of the site option patches described below.

Adjustment of these parameters may prove necessary if the data collected by the GMF/TSSM reveals that excessive MME GEMOREs are being executed during swap file processing. In addition, data collected by the GMF/TSSM can be used to identify the optimum number of swap files which are required, based on the volume of swap processing indicated by the traces collected. GMF/TSSM trace type 2 is generated if an attempt to grow a TSS swap file is denied. The continued occurrence of this type of trace should prompt site personnel to either increase the number of active swap files, depopulate the devices used to hold the swap files, or perform a cold boot to gather fragmented disk space.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000357  OCTAL 000000000004  FOUR SWAP FILES ACTIVE             (#A13)  .MTIMS
000356  OCTAL 000454000000  MINIMUM SWAP FILE GROWTH LINKS     (#A14)  .MTIMS
000355  OCTAL 000000002260  INITIAL SWAP FILE SIZE LLINKS(EACH)(#A15)  .MTIMS
```

Computation of the optimal number and size of the TSS swap files is not exact, and involves consideration of the following factors from the TEARS reports:

- Number of swaps being performed
- Number of force swaps (discussed later in this section)

- o Average subsystem sizes
- o Duration of swaps
- o Configuration limitations on disk and core space
- o Number of concurrent TSS users

The minimum interval between considerations to swap user subroutines can be specified through the modification of the following (non-site option) cell of TSS. The value specified will be used to prevent swap file processing from utilizing too large a share of the total processing time available to the executive. Force swaps may occur without reference to this cell, but are controlled by other methods (as described in subsection 14.6.7.3.7).

```
CC      CC    CC                                                     CC
1       8     1                                                      7
              6                                                      3
000172  OCTAL 000000175000 36/1SEC*64000 CLK PLSES SWAP INVL (A.SD3I) .MTIMS
```

The smallest memory wait time before swap logic is considered can be set to prevent swap logic from being invoked for memory allocation purposes before normal subsystem terminations might allow allocation to be completed.

```
CC       CC    CC                                                     CC
1        8     1                                                      7
               6                                                      3
00000166 OCTAL 372000 36/2SEC*64000 CLK PLSES (TSAWT)                 .MTIMS
```

The following cell is included in the TSS executive, but is not recommended for use, as its ramifications cannot be predicted. If non-zero, the contents of this cell will be used to override any other swap processing timers, forcing swap processing to occur with the frequency specified in the timer. If specified, this cell would seem to be a method by which the number of user subsystems in memory could be minimized by actually swapping each subsystem in a state which will permit it. For the optimization of TSS response time, this would not be recommended. Conversely, if this cell were specified to be high, swapping would not occur when needed for the allocation of large or urgent subsystems. Again, this is not recommended for performance.

```
CC      CC    CC                                                     CC
1       8     1                                                      7
              6                                                      3
000212  OCTAL 000000000 36/MIN*60*SEC*64000 CLK PLSES BTWN SWP-A.SD3C .MTIMS
```

Finally, the following site option patch is recommended for all sites. Examination of the operation of the TSS executive by officials of CCTC/C751 has revealed that that the usage of the various TSS swap files is not balanced. The patch shown below will act to balance this usage, and will minimize the contention caused by imbalanced usage, thus lowering the wait times for those TSS users and subsystems which access them.

```
CC     CC    CC                                                  CC
1      8     1                                                   7
             6                                                   3
011365 OCTAL 012134221200  LDX1 NXFIL LAST FILE USED     #BJS    .MTIMS
011367 OCTAL 042315604200  TMI  P1 RESET INDEX           #BJS    .MTIMS
011373 OCTAL 042317603200  TRC  P2 ARE WE DONE           #BJS    .MTIMS
011376 OCTAL 042317710200  TRA  P2 ARE WE DONE           #BJS    .MTIMS
042315 OCTAL 000357721200  P1  LXL1 .TSSF #SWAP FILS     #BJS    .MTIMS
042316 OCTAL 011366710200      TRA  DRM050 LOOP          #BJS    .MTIMS
042317 OCTAL 012134101200  P2  CMPX1 NXFIL DONE YET      #BJS    .MTIMS
042320 OCTAL 011377600200      TRA   DRM055    YES       #BJS    .MTIMS
042321 OCTAL 011366710200      TRA   DRM050    NO        #BJS    .MTIMS
```

14.6.7.3.4 <u>Subsystem Accounting</u>. The GCOS Statistical Collection File (SCF) software contained in the TSS executive can be adjusted to modify the frequency with which TSS statistical records are written. Site personnel should consider the overhead associated with this statistical record, and review with security personnel the need to collect this information. If it is not used, the site option patches described below can be used to disable its collection and recording. When active, a dummy User Status Table (UST) is maintained by TSS for SCF, thus decreasing the amount of UST (and subsystem) memory available to TSS terminal users, and increasing overhead.

```
CC     CC    CC                                                  CC
1      8     1                                                   7
             6                                                   3
000120 OCTAL 000000000060  COLLECT STATS EVERY 60 SECS   (#A11)  .MTIMS
000352 OCTAL 000000000001  TURN ON SUBSYSTEM ACCNTNG     (#A1B)  .MTIMS
```

On a lightly or moderately loaded TSS system, the presence or absence of subsystem accounting will make little difference to performance. On a more heavily loaded system, however, this site option should be disabled by removing one or both of these site option patches, unless the SCF data collected is required by the site.

14.6.7.3.5 <u>UST Memory Management</u>. Site personnel can modify several parameters which will affect the amount of memory which can be used by the TSS executive for User Status Tables (USTs). These USTs control the number of active users which can be processed by the TSS system concurrently. Active users may be DRUN (Deferred command file) sessions or physical or pseudo (at WIN sites) terminals. Memory for UST entries is set aside directly above the executive code, and is separated from the memory used for user subsystems by a movable fence.

If subsystem accounting is active, one UST will be used by the SCF processes, and will not be available for use. It is included, however, in the maximum UST settings described below.

The maximum amount of memory for USTs is computed based on the maximum number of terminals allowed by the executive. This parameter is a settable option, and should be computed by site personnel based on the following factors, as determined by the reports produced by the TEARS system:

- o Maximum core allowances for TSS, from site configuration limitations as determined from site option patches (section 14.6.7.3.7)
- o Average user subsystem sizes and usage rates as determined from TEARS reports (section 14.6.7)
- o DRUN file usage rate which increases the number of USTs required as determined from the usage of DRL T.CFIO (section 14.6.7)
- o Occurrence of abnormal disconnects, timer setting (see below)

The number of TSS sessions (including DRUN sessions) which will be allowed concurrently is set through the following site option patch.

```
CC      CC    CC                                                          CC
1       8     1                                                           7
              6                                                           3
000117  OCTAL 000000000022  18 USTS MAXIMUM INCLUDING DRUN (.TFMAX)     .MTIMS
```

As mentioned above, this parameter will set the maximum number of USTs which can be active within TSS at any one time. One of the factors which will affect this setting is the rate with which users abnormally exit (i.e. other than BYE) from TSS. The executive will continue to hold the UST for an abnormally exited user for a period, and will allow a user to reconnect to this UST to continue a previous session. The time period for holding these 'Reconnect' USTs is settable, and should be determined based on site needs, maximum core limitations, and the rate with which these abnormal disconnects occur, as determined from the reports produced by the TEARS system.

```
CC      CC    CC                                                          CC
1       8     1                                                           7
              6                                                           3
000310  OCTAL 000035230000  36/2MIN*60*64000 CLK PLSES      (#A12)      .MTIMS
```

The other consideration which will influence the number of active USTs involves the USTs which are used by DRUN, or deferred TSS sessions. The number of sessions which will be handled by TSS, as well as the number of active terminals which will prevent DRUN files from being executed can be set. A preferred time for the scheduling of DRUN sessions can be established. The determination of values for these parameters is dependent upon site requirements. The parameters should be set based on:

- o Site requirements and mission objectives
- o Consideration of the active terminal loads faced by TSS weighted by time of day
- o Site limitations

The maximum number of DRUN sessions which are allowed at a site is dependent upon the usage of this feature. It is set through the use of the following site option patch:

```
CC      CC     CC                                                   CC
1       8      1                                                    7
               6                                                    3
000253  OCTAL  000010000004  8 DRUNS MAX, NONE IF 4 USRS ACTIVE (#A12)  .MTIMS
```

DRUN sessions may be scheduled to begin at a specific time of day, or may be scheduled with no start time specified. If no time is scheduled, the preferred time for scheduling DRUN sessions is used. This parameter is set through the following site option patch, and should be computed based on the lightest active terminal load faced by TSS during the computing day, as discovered through the reports produced by the TEARS system. The default, as shown below, is no preference.

```
CC      CC     CC                                                   CC
1       8      1                                                    7
               6                                                    3
000255  OCTAL  000000000000 36/HOUR*60*MIN*60*SEC*64000      (#A1A)     .MTIMS
```

Finally, the maximum amount of processor time which can be used by any one DRUN session can be limited, thus increasing the throughput of these sessions, and limiting their use of the UST space which might otherwise be used by active terminals.

```
CC      CC     CC                                                   CC
1       8      1                                                    7
               6                                                    3
000254  OCTAL  000000000000  36/MIN*60*SEC*64000 CLK PLSES  (#A19)     .MTIMS
```

GMF/TSSM traces record the manipulation of the fence between UST memory and subsystem memory, and give an indication of the amount of space actually being used, as well as a measure of how dynamic that load is. The settings discussed in this section should be made such that they do not prevent TSS service to users, but should be in line with the actual load which is expected. If 500 terminals are configured at a site, but only 100 are ever seen at once, site personnel should base their limitations on the load experienced, not that which is possible based on the configuration.

14.6.7.3.6 Subsystem Memory Management. Of the factors which can be used to improve TSS response time, those described in this section are the most effective. They control the rates with which TSS will perform the functions which require the most time of the executive. Unfortunately, they are the most difficult to both understand and compute.

14.6.7.3.6.1 Subsystem Memory Allocation Processes. Above the fence which separates UST memory from subsystem memory, and below the base address limitation of TSS, core memory is used to contain active user subsystems. In order for a user subsystem to receive attention from the processor, it must be present in this area. TSS maintains its own swap files, and will force subsystems to these files based on the amount of memory available and the load being experienced.

Based on settable parameters, the TSS executive may request that this user subsystem memory be grown or shrunk by settable increments, and within maximum and minimum TSS size limitations. Allocation routines within TSS contain algorithms which determine the rates with which this subsystem memory is allocated, and are programmed to discriminate against large subsystems to a settable degree. A maximum allowable 'wait time' is computed based on the size of an allocation request. This rate changes when the attempted allocation reaches a 'large subsystem fence' size, when it will become much longer. Once a subsystem has waited for memory allocation for a time greater than this allowable 'wait time', it will become urgent. Once an urgent subsystem has been encountered, the allocation processes will limit allocations for any subsystems which are not urgent by reserving an area of memory for urgent subsystems only.

The occurrence of those GMF/TSSM traces which record the 'urgent user' processes should prompt site personnel to consider modification to one or more of the parameters described in this section. The presence of urgent users (i.e. user memory allocations) is indicative of an overloaded TSS system.

After an urgent user has waited for allocation for a settable time increment, force swaps will be performed against subsystems which have been present in core for a separate minimum amount of time. Subsystems which become 'force swapped' are considered first for allocation after the urgent user has been satisfied. If enough core is not present, or cannot be obtained through force swapping, a size increase may be performed by TSS within its maximum size limitations (discussed in the next section).

Ironically, a subsystem which has been urgent longer than a settable parameter will re-enter the allocation process from the beginning, and allocation will return to normal. Urgent user processing lowers TSS ability to service other allocation requests, as well as line service and other executive functions. the purpose of this parameter is to prevent urgent users from completely monopolizing TSS executive processor time on a heavily loaded system.

Finally, a parameter may be specified which, when exceeded, will abort the allocation for a large subsystem, returning the message 'Not Enough Core' to the requesting terminal. The parameter may be used to prevent or limit the number of times that the urgent user coding is invoked by a single allocation request.

14.6.7.3.6.2 Settable Parameters. The following parameters are associated with the subsystem memory management routines of the TSS executive. Few of these parameters have been established as site option patches. The reader is cautioned against gross modification of these default values which is not both preceded and followed by careful analysis and measurement through the GMF/TSSM and TEARS. The reader is further warned that the modification to any one of these parameters may affect the optimum settings for still others. The complexity of the interrelationships between these parameters has not been approached in this report, and should not be underestimated. Default values used in the examples shown in this section can be assumed to be in balance, and will satisfy requirements at most sites.

The most common modification made by sites which are experiencing response time problems, and the one which will yield the best results are the following patches which bias TSS against allocations for large user subsystems. The patch below establishes the size of a 'large' subsystem, and the penalty (in allowable wait time) which is imposed against allocation requests which are larger than this size.

```
CC      CC     CC                                                 CC
1       8      1                                                  7
               6                                                  3
000200  OCTAL  000044000000  .TAMIS LG SS SIZ = 36K       (#AOX)  .MTIMS
000207  OCTAL  000000000004  .TALPP LG SS WAIT MLTPLR     (#AOX)  .MTIMS
```

By default, a large subsystem is defined as being at least 36K words in size, and it will be forced to wait four times as long as a normal subsystem before becoming urgent. In practice, it is recommended (in general) that the definition of a large subsystem be made based on the actual sizes of the subsystems used at a site, computed based on twice the average of their sizes weighted against the usage these subsystems receive. The wait time penalty should be raised if these subsystems become urgent at times other than the maximum loading periods experienced at the site.

The time period before a subsystem memory allocation is aborted can be modified through the following site option patch. This time period should not prevent a subsystem from becoming urgent, and should be larger than the maximum wait time for the largest subsystem at the site. It should include a margin to allow TSS to attempt to process the subsystem, once overdue. Probable successful values for this parameter will be larger than the maximum urgent user processing time (described later in this section) plus the time required to attempt TSS memory expansion. The value should be lower than some multiple of this sum, to prevent an allocation from becoming urgent more than a few times.

```
CC      CC   CC                                                      CC
1       8    1                                                       7
             6                                                       3
000165  OCTAL 000007640000 36/30SEC*64000 DEFAULT MAX WAIT TME (#AOW) .MTIMS
```

The default value should be modified only if the reports produced by TEARS indicate that allocations for large subsystems are becoming urgent several times before being satisfied.

The maximum amount of time which can be used to attempt urgent user allocations can be limited by using the following patch. This tunable parameter, as with all following it in this section, has not been established as a site option patch. The reader is advised to report any modifications made to these cells when reporting TSS incidents for analysis.

```
CC      CC   CC                                                      CC
1       8    1                                                       7
             6                                                       3
000171  OCTAL 000004704000  36/20SEC*64000 CLK PLSES       (A.MFTM)  .MTIMS
```

The maximum allowable wait time (for all but large subsystems) is computed based on a constant function of the size of the subsystem being allocated. The constant used can be modified to either increase or decrease the amount of time which can be spent by a subsystem waiting non-urgent allocation.

```
CC      CC   CC                                                      CC
1       8    1                                                       7
             6                                                       3
000210  OCTAL 0000030  OCT 30 = 1/6 SEC PER K WAIT TIME DVSR (.TASWF) .MTIMS
```

The reader should observe that lowering the value of .TASWF will increase the amount of allowable wait time, while raising it will decrease wait time. Lowering the wait time will raise the occurrence of urgent allocations, while raising it will decrease urgent allocations (but will increase average allocation times allowed before a subsystem becomes urgent). Urgent allocation processes may be required to load large subsystems, or to force processing-bound subsystems to swap.

The minimum core resident time can be specified with the patch shown below. Subsystems which have been resident for less time than the value of this cell will not be considered for force swap. The purpose of this cell is to prevent 'thrashing' or the constant moving of active subsystems between swap files and core. Of particular interest when computing this value is the TEARS reports which describe the time spent by a subsystem waiting for and receiving processor attention. The values of this cell should be high enough to permit subsystems to receive processor attention.

```
CC     CC  CC                                                  CC
1      8   1                                                   7
           6                                                   3
000173 OCTAL 000001161000 36/5SEC*64000 CLK PLSES MEM RES TIM (A.MTQ) .MTIMS
```

The amount of subsystem memory which can be reserved for urgent user allocations can be limited based on a fraction of the total memory available for subsystems. The actual amount reserved will be equal to the size of the allocation for the largest urgent user. Modification to this cell may force TSS to perform a GEMORE of memory prior to attempting allocation of a subsystem. If the urgent allocation (i.e. memory size for the allocation which has been waiting the longest) is larger than the fraction:

TSS Subsystem Memory Area/A.URMD

Then the urgent user fence cannot be established, and a GEMORE will be forced before the allocation can be satisfied.

```
CC     CC  CC                                                  CC
1      8   1                                                   7
           6                                                   3
000211 OCTAL 000000000002  DEC 2 = MAX URG SIZ 50% SS CORE (A.URMD)  .MTIMS
```

The analyses required to compute the value of this cell are based on the average size of the TSS executive, the number of users (thus the location of the UST/SS memory fence, and the amount of subsystem memory available), the average loading of this subsystem memory (yielding the amount likely to be available and the number of allocations which are likely to be pending), and the size of those subsystems which become urgent most often. If the fraction of the size of these subsystems / the available subsystem memory size is greater than 50% of the total (default), then the urgent user memory fence cannot be established, forcing a GEMORE. If the sizes of the urgent user subsystems cannot be lowered, some adjustment is indicated.

Force swaps will be performed to allow an urgent subsystem allocation to be performed within the memory fence established. The minimum amount of time which must expire before TSS will consider GEMOREs for core size increases to satisfy the allocation are settable based on the cell shown below.

```
CC      CC    CC                                                         CC
1       8     1                                                          7
              6                                                          3
000161 OCTAL  0000175000 36/1SEC*64000 CLK PLSES BFR SIZ INC (.TASID)  .MTIMS
```

The minimum interval between attempts to increase the size of TSS because of unsatisfied urgent allocation requests is settable based on the following cell.

```
CC      CC    CC                                                         CC
1       8     1                                                          7
              6                                                          3
000162  OCTAL 000003523000 36/5SEC*6400 CLK PLSES BTWN GMORE (.TASCF)  .MTIMS
```

Based on the analyses described above, it may be desirable to decrease the minimum time which must expire before a size increase can be considered for large urgent allocations. Based on the failure rate for GEMORE requests (as determined from GMF/TSSM traces, and caused by a heavy non-TSS core load), the interval between these requests might also be raised.

The maximum amount of time which will be spent by TSS in the attempt for a single GEMORE request can be controlled, to limit the performance degradation faced by TSS users in situations where growth is not possible. Reasons why this might occur include programs marked 'dead' directly above TSS in core, and are usually transitory in nature. Should response times be affected by the length of time spent by TSS attempting to GEMORE memory, this cell might be changed to a lower value.

```
CC      CC    CC                                                         CC
1       8     1                                                          7
              6                                                          3
000163  OCTAL 000016514000 36/60*SEC*6400 CLK PLSES GMOR LMT (.TATMC)  .MTIMS
```

The interval between GEMORE attempts to expand TSS memory can be controlled, preventing TSS from attempting to grow too fast, and limiting the time spent in considering size changes. The following cell, in conjunction with .TAMRI (discussed below) will limit the rate with which TSS will change size.

```
CC      CC    CC                                                         CC
1       8     1                                                          7
              6                                                          3
000164 OCTAL  002342000 36/10*SEC*64000 CLK PLSES GMR LMT ATT(.TAGMI)  .MTIMS
```

As with size increase considerations, the intervals between size reduction considerations and attempts can also be controlled. If the intervals described above for size increases are shortened, permitting rapid acquisition of space by TSS, then consideration should also be given to making similar changes to the cells described below, permitting rapid return of that space to the system possible also. If corresponding changes are not made to this logic, TSS will monopolize core memory longer than required to satisfy its subsystem requirements.

The minimum interval between size reduction considerations can be controlled based on the usage of the TSS swap files. If the swap file usage exceeds the value of the cell shown below, consideration of size reductions will be prevented. Situations may arise, and may be reported by TEARS, where TSS appears to be too large, but swap file usage (as indicated by GMF/TSSM traces) will show that numerous swap-ins are in progress.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000206 OCTAL  0000000113  DEC 75% SWP FIL UTL PREVNT RDCTN (.TAPMR)   .MTIMS
```

If swap file usage is not in excess of .TAMPR, the cell shown below will limit the number of times that a TSS size reduction will be considered.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000157  OCTAL 0035230000 36/2MIN*60*6400 CLK PLSES BTWN RDCTN .TAMRI  .MTIMS
```

After the need for a TSS size reduction is identified, a settable delay will ensue before it is actually performed. This delay will prevent rapid core releases from being scheduled (which would likely be followed by rapid core requests). This damper acts to synchronize the clocks which limit considerations of TSS size reductions and size increases.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000160  OCTAL 0002342000 36/10SEC*64000 REDUCTION DELAY    (.TATMD)   .MTIMS
```

According to documentation and program comments, successful values for this cell should be less than the sum of .TATMC (maximum size change wait interval) plus two times .TAGMI (minimum interval between GEMORE requests).

The final settable parameter which relates to the user subsystem memory area concerns not user subsystems but rather VIP terminal buffers. The Buffered Terminal Output System (BTOS) will allocate Extended Memory Buffers (EMBs) in the subsystem memory area to hold terminal output prior to its being transmitted to terminals. The site option patch shown below can be used to limit the number of VIP terminals which can be active (on-line) at any one time, and thus the number of EMBs which might be active.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000116  OCTAL 0000000000 18/NUMB OF VIP 760,18/0           (.T760)    .MTIMS
```

14.6.7.3.7 TSS Executive Processing. The site option patches described in this section relate to the TSS executive performance, and will affect the time required to perform many of the functions related to TSS response.

TSS maintains its own Peripheral Allocation Tables (PATs) which define the location and size of the files accessed by the executive, as well as those accessed by each time sharing user active in the system. If analysis of the GMF/TSSM data collected during a monitoring session reveals the presence of trace type 3s, then the amount of space available to TSS to maintain these PATs is insufficient. PAT tables are maintained by TSS in its Slave Service Areas (SSAs). The number of SSAs required is computed by TSS based on the expected size of a PAT entry times the number of entries expected to service the maximum number of TSS users which are allowed. There are two ways in which this computation may need to be overridden. First, the number (or size) of the files used by TSS users may be greater than that expected by TSS. Second, disk space fragmentation might expand the size of the PAT entries required to define the physical locations of the files accessed by TSS users. When GMF/TSSM trace type 3s are encountered, the number of SSAs reserved by TSS (and hence, the space available for PAT entries) may be modified through the use of the following patch.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000136  OCTAL 000000000000 0=COMPUTE BASED ON MAX USERS    .TSSA      .MTIMS
```

The minimum interval between entries to the Line Service state (section 14.6.7.1) to process normal terminal I/O requests, terminations, and status checks is controlled by the following non-site option parameter. If analysis of the GMF/TSSM data tapes produced during monitoring sessions reveal that the line service state is being entered without any need, then this interval may be raised. If analysis reveals that the volume of work faced by this remote I/O module is high, the interval may be lowered to improve response.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
000167  OCTAL 000000076400 36/.5SEC*64000 LINE SVC ENTRY   (.TLTIM)   .MTIMS
```

Periodic service performed by the TSS executive Line Service module is performed according to the setting of the parameter shown below. Periodic service functions include noticing of terminal disconnects and new log-ons. While not a significant element of terminal response time, this factor might be raised if analysis reveals that the number of terminal connects and disconnects is very low. The possible savings in terms of processing time are not high, but might be worth considering on a heavily loaded system.

```
CC      CC   CC                                                        CC
1       8    1                                                         7
             6                                                         3
000170  OCTAL 0000567000 36/3SEC*64000 CLK PLSES PERIOD SVC NTRY(#A10) .MTIMS
```

Finally, the maximum size, minimum size, amount to obtain, and amount to release for core storage can be specified by the following site option patch. Minimum size requirements are those which will be sufficient to satisfy the load experienced by TSS during periods of off-peak use. The difference between heavy and light use must be defined by each site. During periods of low (off-peak) use, TSS should not be forced to GEMORE core. The amount of core required by TSS at the heaviest point in its processing day should be measured through the GMF/TSSM and used as the maximum size. Increments to GEMORE and release should be developed based on the frequency with which these functions are performed. Consideration should be given to the sizes of the subsystems which most commonly cause size increases to occur. The frequency with which these subsystems are used and the amount of core required to run them will indicate the amount of core which could be GEMOREd by the TSS executive. In general, GEMORE and release core increments are made based on the weighted average core requirements which cause size increases to be considered. The parameter for core release should be an even multiple of that used for GEMORE, for obvious reasons.

```
CC      CC   CC                                                        CC
1       8    1                                                         7
             6                                                         3
000176  OCTAL 000000000144 .TAMMS MAX TSS CORE IN K   DEC 100  (#A1C)  .MTIMS
000177  OCTAL 170000000000 .TASMS MIN TSS CORE IN K   DEC 60   (#A1C)  .MTIMS
000201  OCTAL 000000000024 .TAMII MEM ADD INCREMENT   DEC 20   (#A1C)  .MTIMS
000202  OCTAL 024000000000 .TASRI MEM REL INCREMENT   DEC 10   (#A1C)  .MTIMS
```

The computation of the values of these parameters is made as follows. .TAMMS is the octal number of K words which limits the maximum size of the executive. Bits 0-17 of .TASMS control the minimum size of the TSS executive in words (not K words). The altered value of this variable is equivalent to 45K, and is considered the minimum reasonable number. .TAMII is the memory addition increment in K words. .TASRI, bits 0-17 are the number of words (not K words) which should be used for memory releases.

14.6.7.3.8 Subdispatch Queue Processing. With WWMCCS software release W7.2B (commercial 4JS), a change was made to the processes used to grant processor attention to active user subsystems. The system dispatching module, .MDISP, was modified to process a subdispatch queue for both TSS and Transaction Driven System (TDS). The purpose of these subdispatch queues is to grant privleged attention to those portions of the system which are interactive, or which will affect the response times experienced by terminal users.

There are several parameters which can be adjusted to affect the degree to which these subdispatch queues will receive privleged attention from GCOS. An understanding of the processes involved is important before any adjustment is attempted.

The first of the three subdispatch queues used by TSS is the Available queue, which is a linked list of all queue entries which are not currently linked into either the Ready or Fault queues. When a subsystem has been loaded into the Subsystem memory area, an entry is obtained from this available queue, loaded with the address of the subsystem, and linked into the Ready queue. Once placed in this Ready queue, no further GMF/TSSM traces for this subsystem (and hence, the user) will be encountered until the subsystem has received attention from the processor.

.MDISP, during its processing, will remove an entry from the Sub-dispatch Ready queue when selecting a new slave process for dispatch. Tunable parameters exist which can be used to adjust how often this selection will be made from the subdispatch queue as opposed to normal batch slaves.

Once a subsystem has been dispatched to, some fault will occur to end the dispatch. For most subsystems, this will be a DRL fault signifying either an I/O request or subsystem termination. In certain cases, the fault may be a timer runout, but this will occur for only a small percentage of processor-bound subsystems. When this fault occurs, .MDISP will place the subsystem address in an entry to the Subdispatch Fault queue. During executive processing, TSS will retrieve this entry from the Fault queue, and act to resolve the fault encountered. Subsystems which have received timer runout faults will be placed back in the Ready queue for further processor attention. DRL faults must first be processed before the subsystem can become eligible for processor attention again.

Investigation has shown that, in certain circumstances, subsystems which are heavily processor-bound can degrade system performance by monopolizing the available processor attention. While this will not greatly affect other user subsystems (Since the Ready queue is processed sequentially), it will limit the attention which will be given to other batch jobs, because of the special attention given to subdispatches. Log number #BNF has been developed for release W7.3.0 to minimize this potential problem by limiting the number of consecutive subdispatches (without an intervening I/O or other fault) a subsystem can receive before a delay is forced for that subsystem. Before its release, this patch may become a site option.

Current feelings are that there are other tunable parameters which can also be used to limit this problem.

Within TSS, the following tunable parameters can be used to affect sub-dispatch queue processing. The patch below can be used to affect the duration of a processor time-slice which will be granted to a subsystem which has been placed in the Ready queue.

```
CC      CC    CC                                                        CC
1       8     1                                                         7
              6                                                         3
002630  OCTAL 000000002400    36/20*64000  CLK PLSES        (.QQTM)    .MTIMS
```

Also within TSS, the patch below can be used to affect the threshold of Ready queue entries which will force the executive to perform certain actions during its own dispatch.

```
CC      CC    CC                                                        CC
1       8     1                                                         7
              6                                                         3
002641  OCTAL 000012000004         RDY QUE LMT, THRESHOLD   (.QRCL)    .MTIMS
```

.QRCL is used by TSS in the following manner during an executive dispatch:

- o IF .QRCL (18-35) is greater than the number of entries which are in the Ready queue, TSS will relinquish to allow subsystem dispatching.
- o IF .QRCL (18-35) is less than the number of entries waiting in the Ready queue, and if both are lower than the limit in .QRCL (0-17), the executive will first process Fault queue entries before relinquishing to allow subsystem dispatching.
- o IF the number of entries in the Ready queue is greater than .QRCL (0-17), then the executive will enter the Line Processing state to handle any outstanding remote I/O requests.

Modification to .QRCL will affect the speed with which TSS will perform its executive functions, thus affecting subsystem allocation and fault processing. By raising the value of .QRCL (0-17), the number of times that TSS will attempt to service both outstanding faults and remote I/O requests will decrease. By raising the value of .QRCL (18-35), subsystem processor attention will be increased, but the number of times that the Sub-dispatch Fault queue will be processed will be decreased. It is suggested that, before attempting to adjust .QRCL, site personnel first attempt to improve response time by adjusting the length of both an executive dispatch (to decrease fault queue wait time by allowing TSS to process all outstanding faults during a single dispatch), as well as .QQTM (to allow subsystems a longer processor time slice).

Within the system dispatching module, .MDISP, the decision as to which process will receive the next dispatch is made based on two dampers. Altered values will allow six subdispatches to occur for every two normal batch dispatches. The patch shown below can be used to affect this balance.

```
CC      CC    CC                                                      CC
1       8     1                                                       7
              6                                                       3
001305  OCTAL 000000000006 QDCT SUB-DSP COUNTER RFSH                  .MDISP
001307  OCTAL 000000000002 QNDR BATCH COUNTER RFSH                    .MDISP
```

Consecutive dispatches are made by .MDISP until one counter has been emptied, or until no further (subdispatch or batch) processes are waiting, before the other type of process is given attention. Thus, six sub-dispatches will be made, followed by two normal batch dispatches (One of will probably be to the TSS executive if Priority 'B' processing is active), given the default values shown. As with the threshold to .MTIMS shown above, it is suggested that site personnel first attempt to affect performance by varing the length of a processor dispatch before attempting to modify the ways in which dispatches are granted.

14.6.7.4 Identification of TSS Response Time Problems. The performance information gathered by the GMF/TSSM can be used by site personnel to identify the frequency of occurrence of response time problems. Delays which occur can be caused by any number of factors. These factors can be either related or unrelated to user actions. This section describes the ways in which the reports produced by TEARS can be used to identify the occurrence of these response time problems.

14.6.7.4.1 General. Normally, site personnel will first encounter a response time problem as the result of user complaints. Once identified by the user, the GMF/TSSM can be executed over a period of days (or weeks) before the particular problem noticed by users is encountered again. Typically, a user is unable to reproduce a problem at will, as most are related to either system loads or some other uncontrollable circumstance. If site personnel have incorporated the suggestions contained in Section 14.6.7.3 of this report, many of the problems which are typically encountered can be minimized or prevented.

In order to identify the existence of a problem, (if, in fact, any problem exists) site personnel must first attempt to categorize the nature of the complaint. The answers to the following questions are of paramount importance in this categorization:

- o What was the user doing? What subsystem was he attempting to run, how did the delay manifest itself?

- o How long was the user delayed?

- o How many users were affected? Were they all trying to do the same (type of) thing?

- o What USERID was being used? What terminal ID (very important in the case of shared accounts)?

- o During what time of day did (does) the problem occur ?

- o Can the problem be recreated, or forced to occur at will?

Without answers to at least most of these questions, user problems with terminal response time will be difficult at best to even identify. The volume of information collected by the GMF/TSSM will almost preclude detailed analysis of all of the trace information collected over a normal processing day.

If the nature of the complaint or problem experienced seems to indicate a continuing or worsening situation, site personnel should execute the GMF/TSSM to collect trace information. The GMF/TSSM collector should be executed several times, under varing loads, and during different times of day, over several days before attempting to reduce and analyse the collected data. In this way, the site analyst can be better assured that he h s collected sufficient information for resolution. Whether attempting to analyse a specific problem, or to optimize performance for all us s, the collection of trace data over a period, and under varying loads, s important to allow for an understanding of the usage profile at a specific site.

Once sufficient GMF/TSSM trace data has been collected, Site personnel should begin the process of response problem identification with the TEARS RESPONSE phase of data reduction. The reader is referred to the GMF Users Manual for direction concerning the execution of TEARS.

14.6.7.4.2 Analysis of TEARS RESPONSE Phase Reports. In this section, each of the reports produced by the RESPONSE phase of the TEARS data reduction system is examined, and some direction as to the type of problem each can indicate is described.

14.6.7.4.2.1 Timesharing Reduction Event Log. The first report produced by the RESPONSE phase of the TEARS system is the event log. Figure 14-15 illustrates this report. The importance of the log lies in i's use to identify the location within the data tapes of significant events during a tracing session. As described in section 14.6.7.3.1, it would be much easier if only a portion of the trace information had to be analysed. This report should be used to identify the portion by both record number and time of day. For example, if a TRAX-EXEC user on terminal DA reports a problem, then only that portion of the GMF/TSSM data tape from record number 737626 to 1051854 would be of interest. This event log should be retained as a permanent index to any GMF/TSSM data tape which is used for further analysis.

```
                    TIME SHARING REDUCTION EVENT LOG

0236:21.24 ***** MAIN INITIALIZED *****  60.00 SEC. PLOT INTERVAL *****
0237:19.55       0 *** TIMEFRAME ********* BEGIN *****
0237:19.55       1= 90=    =     = ***** TRACE ON *****
0237:19.55       3= 90=    =     = *****   2 USERS *****
0239:17.95    3707= 12=<!
```

```
0240:23.97    4196= 32=<TM>=NEW = LOGON    B29IDPXD
0242:24.94   19672= 32=<TN>=NEW = LOGON    B29IDPXD
0243:30.65   37267= 32=<TO>=NEW = LOGON    B29IDPXD
0245:19.80   81107= 32=<TP>=NEW = LOGON    B29IDPXD
0251:43.36  244048= 22=<TN>=CARD=   LINE SET FOR RECONNECT
0252:12.41  256163= 32=<TN>=NEW = LOGON    B29IDPXD
0254:39.86  314658= 12=<TN>=TERM= LOGOFF   B29IDPXD     ON AT  242.25
0259:36.77  424135= 32=<TO>=NEW = LOGON    B29IDPXD
0301:45.44  470646= 32=<T2>=NEW = LOGON    B29IDPXD
0304:36.85  529279= 71=    =     = ** UST AREA INCREASE.    8 USERS.
0304:36.93  529313= 32=<T3>=NEW = LOGON    B29IDPXD
0305:46.83  554224= 22=<TL>=NONE=   LINE SET FOR RECONNECT
0306:08.18  561659= 32=<TL>=NEW = LOGON    TRAX-EXEC
0308:44.88  613808= 12=<TL>=TERM= LOGOFF   MASA         ON AT  236.30
0309:51.25  635563= 32=<TD>=NEW = LOGON    B29IDPXD
0315:04.87  737626= 32=<DA>=NEW = LOGON    TRAX-EXEC
0315:07.28  738495= 71=    =     = ** UST AREA INCREASE.   10 USERS.
0327:50.38  996338= 22=<DA>=CMDL=   LINE SET FOR RECONNECT
0330:46.92 1051854= 12=<DA>=TERM= LOGOFF   TRAX-EXEC    ON AT  315.04
0331:46.60 1071854= 32=<DA>=NEW = LOGON    6
0336:56.37 1173949 ***** REDUCTION ENDED * T# 23***** STACK ERRORS=       0
```

Figure 14-15. **Time Sharing Reduction Event Log Report**

14.6.7.4.2.2 Response Times for all Users. This is the first plot report produced by TEARS. The report provides the analyst with two very important types of information. First, an indication of the average response times being experienced is directly displayed. Second, anomalies to this average can be first discovered and noted. Figure 14-16 contains an example of this type of report. Note the average response times which have been marked with a pen. Compare how this average rises with the number of active terminals (left hand column). The LINE column displays the user who experienced the longest delay during the time interval shown on each line of the report.

The example shown in figure 14-16 shows that between 3:00 and 3:01, the minimum response time experienced suddenly jumped to over 30 seconds. This type of jump or difference should be noted, and the time of day should be compared with other reports.

This first plot cannot reveal any more than the fact that the response time was different during some interval. No indication of the source of this difference can be found here.

The user of TEARS is advised that the trace information shown at the beginning and ending of the tracing period is skewed due to the fact that incomplete intervals are shown. For example, the GMF/TSSM may record the end of a response interval just after the beginning of the tracing period, but this value will be skewed downward because TEARS must assume that the response interval began just before the beginning of trace data collection. A similar situation exists at the end of a tracing period, when the beginning of an interval is noted, but tracing is shut off before the response interval completes.

The validity of this report, however, is comprimized by the fact that it actually displays wait times, not response times. In the example shown, terminal TM is executing a BASIC program which is trapped in an endless loop. Thus, no responses are given to terminal TM at all between 2:41 and 3:33, but the response time experienced by that terminal seems to be steadily degrading. Because of this significant difference in definition, terminal TM will mask any response problems for terminals which were truly interactive during this interval.

This problem is shared by the next two RESPONSE TIMES reports described below.

14.6.7.4.2.3 Response Times For User Not Requesting More Core. The example shown in figure 14-17 is a subset of the information shown on the All USERS report described above. The difference is the fact that this report shows response times for users who have not requested core during the time frames shown. What this means is that they are within a subsystem, and have not been swapped by TSS. Again, an anomaly exists at 3:01, with terminal TQ experiencing the most delay.

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

RESPONSE TIME FOR ALL USERS

PLOT- 1

-= MIN /= AVG *= MAX

USERS	TIME	WAITS IN INTERVAL (SEC) <1	<2	<4	<8	<16	<30	>30	LINE
2	238:19.56	.	.	.	.	.	.	1	<!!>
1	239:19.56	9	1						<TL>
1	240:19.56	6	.	.	.	1			<TL>
2	241:19.56	16							<TM>
2	242:19.56	12	.	..	.	.	.	1	<TM>
3	243:19.56	13	.	.	.	.	.	1	<TM>
4	244:19.56	13	.	.	.	.	.	2	<TM>
4	245:19.56	4	.	.	.	.	.	2	<TM>
5	246:19.56	14	1	.	.	.	.	2	<TM>
5	247:19.56	40	3	.	.	.	.	2	<TM>
5	248:19.56	41	.	1	.	.	.	2	<TM>
5	249:19.56	35	2	.	.	.	.	2	<TM>
5	250:19.56	37	.	.	.	.	.	2	<TM>
5	251:19.56	48	5	.	.	.	.	2	<TM>
6	252:19.56	12	.	.	.	.	.	2	<TM>
6	253:19.56	19	1	.	.	.	.	2	<TM>
6	254:19.56	8	.	1	1	.	.	3	<TM>
5	255:19.56	7	.	.	3	1	.	3	<TM>
5	256:19.56	2	.	.	2	.	.	3	<TM>
5	257:19.56	2	.	.	2	.	.	3	<TM>
5	258:19.56	5	.	.	3	.	.	3	<TM>
5	259:19.56	3	.	.	2	.	.	3	<TM>
6	300:19.56	15	.	.	.	.	1	3	<TM>
6	301:19.56	.	.	.	.	.	.	5	<TM>
7	302:19.56	2	.	.	.	.	.	5	<TM>
7	303:19.56	2	.	.	1	1	.	5	<TM>
7	304:19.56	9	.	.	.	.	.	6	<TM>
8	305:19.56	6	2	.	.	.	.	6	<TM>
9	306:19.56	12	2	.	.	.	.	6	<TM>
9	307:19.56	7	.	.	.	.	.	7	<TM>
9	308:19.56	6	.	1	.	.	.	7	<TM>
8	309:19.56	7	.	.	.	.	.	7	<TM>
9	310:19.56	12	1	1	1	.	.	7	<TM>
9	311:19.56	11	1	2	.	.	.	7	<TM>
9	312:19.56	7	.	.	.	.	.	8	<TM>
9	313:19.56	10	.	.	.	.	.	8	<TM>
9	314:19.56	8	.	.	.	1	.	7	<TM>
10	315:19.56	19	1	.	3	.	.	7	<TM>
10	316:19.56	27	1	4	2	.	.	7	<TM>
10	317:19.56	37	2	2	5	.	.	7	<TM>
10	318:19.56	28	1	5	3	1	.	7	<TM>
10	319:19.56	14	.	1	4	.	.	8	<TM>
10	320:19.56	13	1	3	3	.	.	8	<TM>

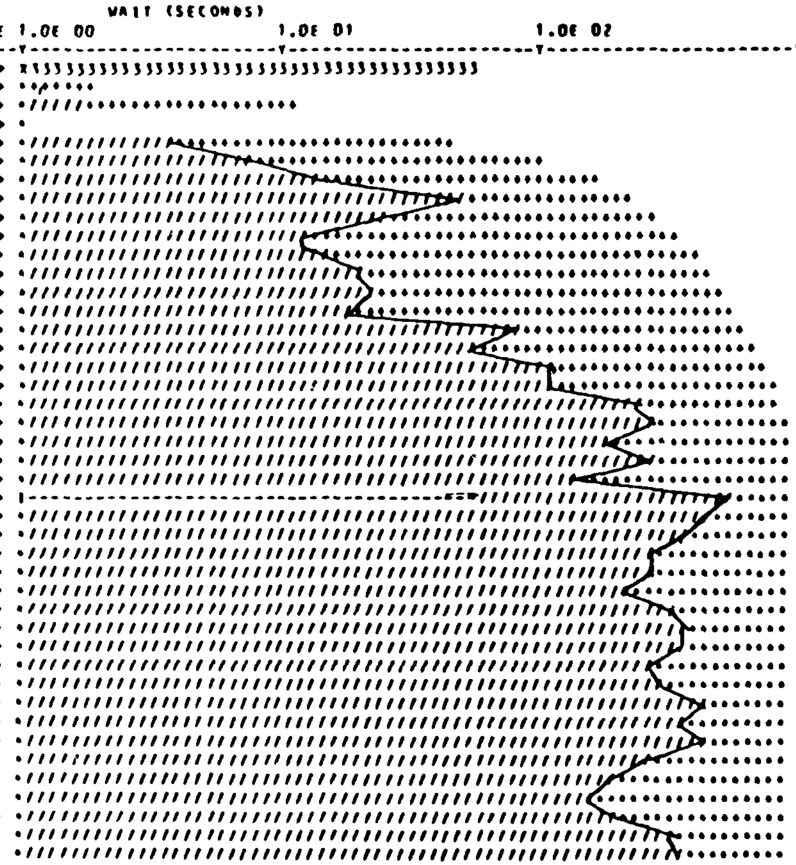

Figure 14-16. Response Times For All Users Report (Part 1 of 2)

USERS	TIME	<1	<2	<4	<8	<16	<30	>30	LINE
10	321:19.56	17	1	2	3	.	.	8	<TM>
10	322:19.56	13	.	1	5	.	.	8	<TM>
10	323:19.56	19	1	1	2	1	1	7	<TM>
10	324:19.56	9	.	.	3	.	.	8	<TM>
10	325:19.56	13	.	1	2	1	.	8	<TM>
10	326:19.56	7	.	2	.	.	.	8	<TM>
10	327:19.56	8	.	2	.	.	.	9	<TM>
10	328:19.56	10	.	.	1	.	.	8	<TM>
10	329:19.56	11	.	1	.	.	.	7	<TM>
10	330:19.56	13	.	.	.	.	.	7	<TM>
9	331:19.56	15	.	.	.	.	.	7	<TM>
10	332:19.56	34	3	1	.	1	.	7	<TM>
10	333:19.56	56	3	.	.	.	.	4	<TM>
10	334:19.56	57	.	.	.	.	1	1	<TQ>
10	335:19.56	51	.	2	2	1	.	4	<TQ>
10	336:19.56	30	2	.	3	.	.	6	<TQ>
10	336:56.37	15	.	.	5	.	.	6	<TQ>

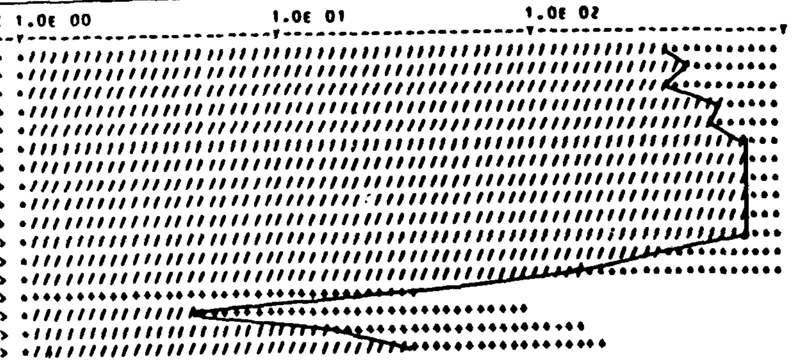

CURVE	MINIMUM	MAXIMUM
MIN	4.2300E-04	6.0540E 01
AVG	1.0117E-01	7.3386E 02
MAX	4.1699E-01	3.0825E 03

Figure 14-16. (Part 2 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

RESPONSE TIME FOR USERS NOT REQUESTING MORE CORE PLOT- 2

-= MIN /= AVG += MAX

USERS	TIME	WAITS IN INTERVAL (SEC) <1	<2	<4	<8	<16	<30	>30	LINE
2	238:19.56	.	.	.	.	.	.	1	<!!>
1	239:19.56	5							<TL>
1	240:19.56	5							<TL>
2	241:19.56	11							<TM>
2	242:19.56	10	.	.	.	.	.	1	<TM>
3	243:19.56	8	.	.	.	.	.	1	<TM>
4	244:19.56	7	.	.	.	.	.	2	<TM>
4	245:19.56	4	.	.	.	.	.	2	<TM>
5	246:19.56	8	.	.	.	.	.	2	<TM>
5	247:19.56	28	.	.	.	.	.	2	<TM>
5	248:19.56	24	.	.	.	.	.	2	<TM>
5	249:19.56	24	.	.	.	.	.	2	<TM>
5	250:19.56	20	.	.	.	.	.	2	<TM>
5	251:19.56	26	.	.	.	.	.	2	<TM>
6	252:19.56	9	.	.	.	.	.	2	<TM>
6	253:19.56	17	1	.	.	.	.	2	<TM>
6	254:19.56	8							<TL>
5	255:19.56	6							<TL>
5	256:19.56	2							<TL>
5	257:19.56	2							<TL>
5	258:19.56	5							<TL>
5	259:19.56	3							<TL>
6	300:19.56	9	.	.	.	.	1		<TO>
5	301:19.56	.	.	.	.	.	.	2	<TO>
7	302:19.56	.	.	.	.	.	.	2	<TO>
7	303:19.56	1							<T2>
7	304:19.56	5	.	.	.	.	.	1	<T2>
8	305:19.56	4							<T3>
9	306:19.56	10							<T3>
9	307:19.56	7							<T3>
9	308:19.56	6							<T3>
8	309:19.56	7							<T3>
9	310:19.56	9							<TD>
9	311:19.56	11	.	.	.	.	.	1	<TD>
9	312:19.56	7	.	.	.	.	.	1	<TD>
9	313:19.56	8	.	.	.	.	.	1	<TD>
9	314:19.56	8							<T3>
10	315:19.56	14							<T3>
10	316:19.56	19							<T3>
10	317:19.56	28	1						<DA>
10	318:19.56	23							<TL>
10	319:19.56	14							<T3>
10	320:19.56	12							<T3>

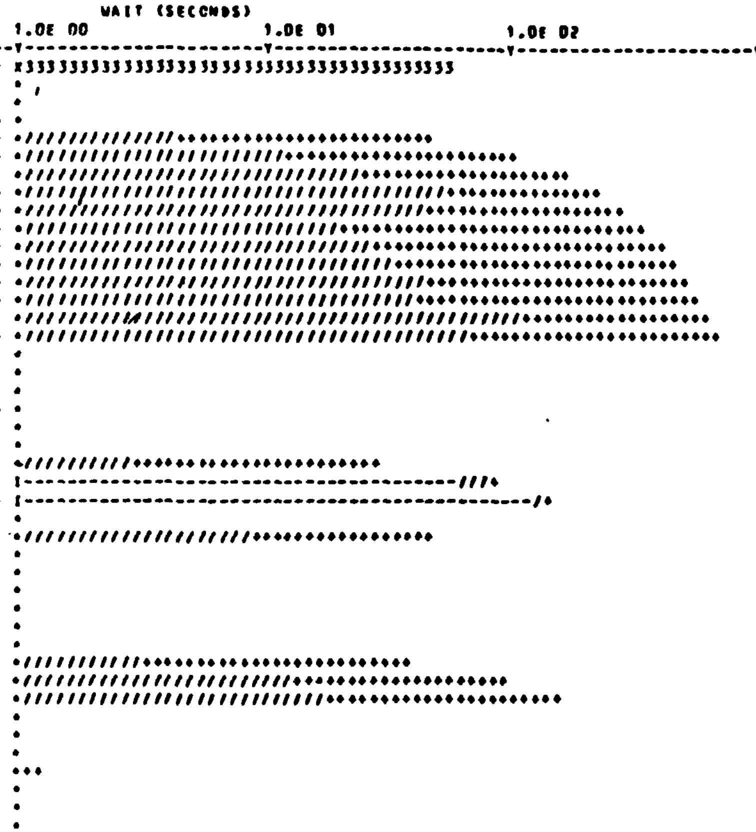

Figure 14-17. **Response Times For Users Not Requesting More Core Report (Part 1 of 2)**

```
USERS   TIME          <1   <2   <4   <8  <16  <30  >30  LINE 1.0E 00                   1.0E 01                   1.0E 02
---------------------------------------------------------------Y-------------------------Y-------------------------Y-------------------------Y
 10    321:19.56      17                                <T3> *
 10    322:19.56      13                                <T3> *
 10    323:19.56      16                                <DA> *
 10    324:19.56       9                                <T3> *
 10    325:19.56      12                                <T3> *
 10    326:19.56       9                                <T3> *
 10    327:19.56       7                                <T3> *
 10    328:19.56      10                                <DA> *
 10    329:19.56      10                                <T3> *
 10    330:19.56      12                                <TL> *
  9    331:19.56      14                                <T3> *
 10    332:19.56      24                                <DA> *
 10    333:19.56      38                                <T3> *
 10    334:19.56      37                                <DA> *
 10    335:19.56      35    .    .    .    1            <T2> **************************
 10    336:19.56      22                                <TN> *
 10    336:56.37      14                                <T3> * ,

   CURVE        MINIMUM         MAXIMUM
   MIN        4.2300E-04      1.2054E 02
   AVG        7.2600E-04      1.3406E 02
   MAX        7.2800E-04      7.0752E 02
```

Figure 14-17. (Part 2 of 2)

The information in this report is best used in comparison to the next report below for users who do have outstanding core requests. The magnitude of the differences in response times shown during common times for these reports is an indication of the times required to perform memory allocation and loading. Comparison of these two reports to the first ALL USERS report should (but does not well) indicate the effect of memory allocation, user subsystem swapping, and subsystem I/O on response times.

14.6.7.4.2.4 Response Times For Users With Core Request During Line Idle. This report is show in figure 14-18. This third plot is the second subset of the overall response time indications given by the ALL USERS response times report described in section 14.6.7.4.2.2. This subset (of the overall response times shown in the first RESPONSE report) is based on those users (subsystems) which required a memory allocation during the response time interval shown. A memory allocation is made for users who are either invoking a new subsystem (i.e. CARD N) or whose subsystem has been swapped and is now being unswapped. The response times noted in this report will always be higher than those in the 'NO CORE REQUEST' plot, and may in fact, reveal the sources of anomalies noted in the 'ALL USERS' plot.

Note the anomaly discovered earlier at time 3:01, which also appears in this plot. It seems that the major element (in the 'ALL USERS' plot) for this delay concerned terminal TM, who was attempting either a subsystem load or an unswap. Indeed, in the example shown, terminal TM is shown each time the minimum response time is high. Either the subsystem being used by this terminal is very large (as was the case) or terminal TM is processor bound thus becoming a swap candidate (also the case) or terminal TM is attempting to load a very large subsystem. If this were an initial subsystem load, some indication could be gained from several of the reports described below.

14.6.7.4.2.5 Total Time In Subdispatch Queue Report. The SUBDISPATCH reports described in the following sections can be used to identify the nature of the work being performed to TSS subsystems during the tracing period. Of particular interest in the report shown in figure 14-19 (which is repeated in each of the next two reports) is the column labeled TROUT. TROUT is a percentage of the subdispatches which occurred during the time frame just ended which ended with a timer runout fault (as opposed to a DRL or some other fault). In the example shown, it appears that all users are processor bound, and could be greatly affected by some increase in the amount of processor time per sub-dispatch (see section 14.6.7.3.9). The BSY column provides a vague measure of how busy the line service module of TSS has become during the interval. The example shown indicates that the subsystems in execution do not perform much I/O. CPU is a measure of the queue processing overhead faced by .MDISP, but since this cannot be affected by site personnel, its purpose remains unclear.

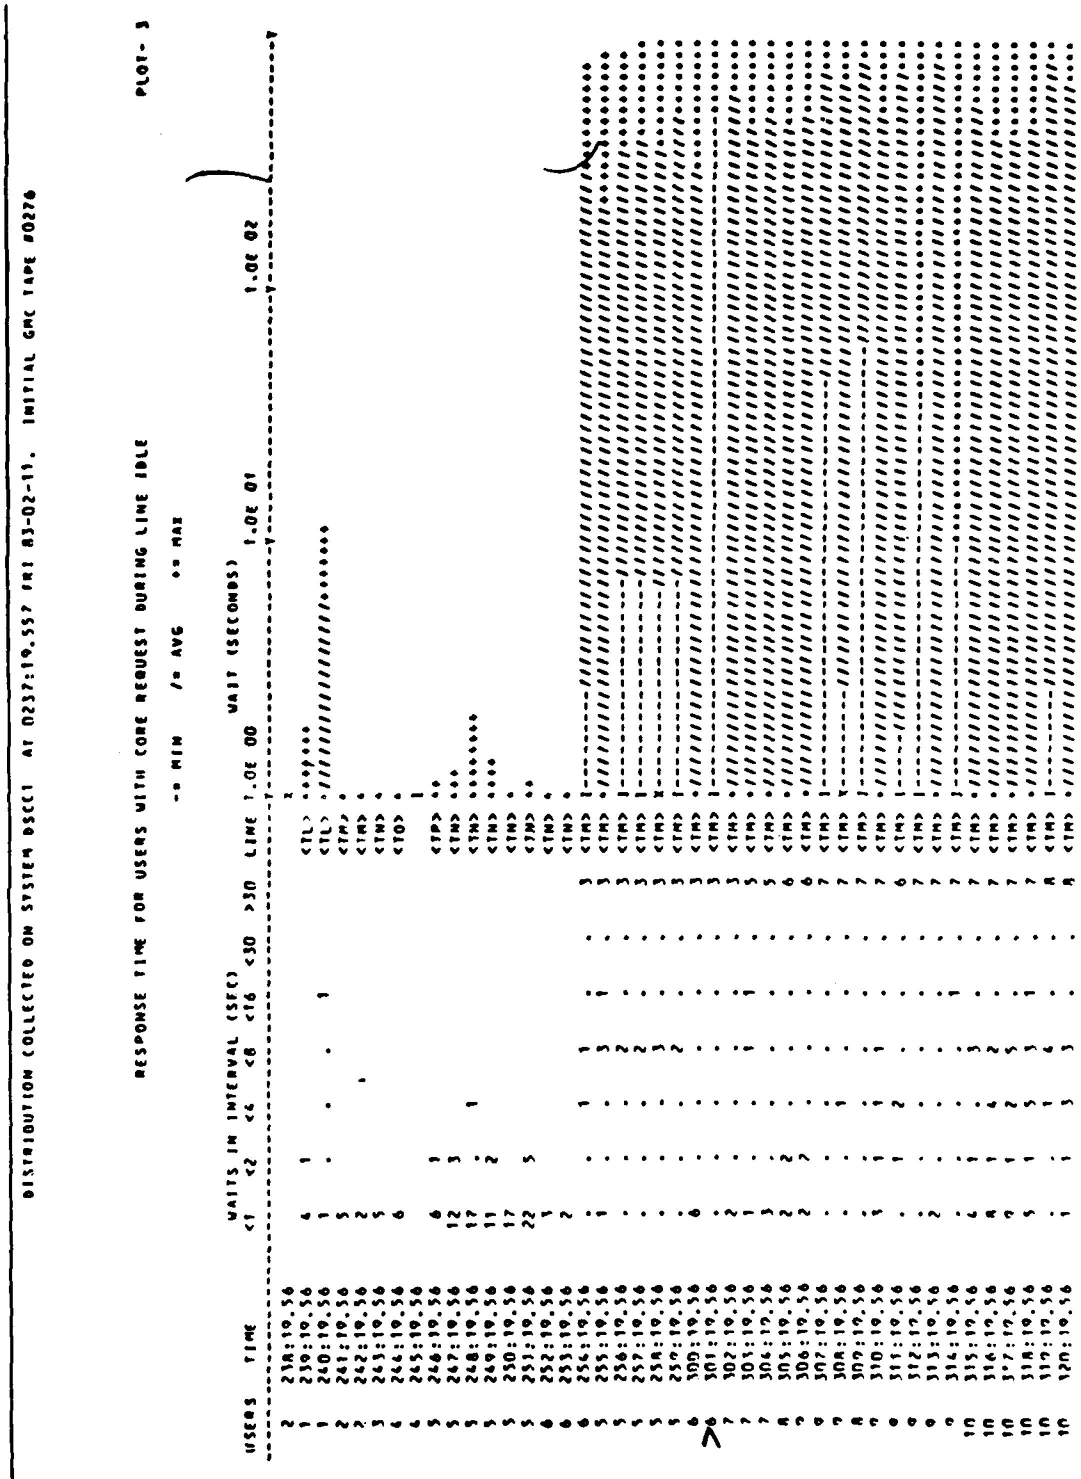

Figure 14-18. Response Times For Users With Core Request During Line Idle Report (Part 1 of 2)

USERS	TIME	<1	<2	<4	<8	<16	<30	>30	LINE
10	321:19.56	.	1	2	3	.	.	8	<TM>
10	322:19.56	.	.	1	5	.	.	8	<TM>
10	323:19.56	3	1	1	2	1	1	7	<TM>
10	324:19.56	.	.	.	3	.	.	8	<TM>
10	325:19.56	1	.	1	2	1	.	8	<TM>
10	326:19.56	.	.	2	.	.	.	8	<TM>
10	327:19.56	1	.	2	.	.	.	9	<TM>
10	328:19.56	.	.	.	1	.	.	8	<TM>
10	329:19.56	1	.	1	.	.	.	7	<TM>
10	330:19.56	1	.	.	.	.	.	7	<TM>
9	331:19.56	1	.	.	.	.	.	7	<TM>
10	332:19.56	10	3	1	.	1	.	7	<TM>
10	333:19.56	18	3	.	.	.	.	4	<TM>
10	334:19.56	22	.	.	.	.	1	1	<TO>
10	335:19.56	16	.	2	2	.	.	4	<TO>
10	336:19.56	8	2	.	3	.	.	6	<TO>
10	336:56.37	1	.	.	5	.	.	6	<TO>

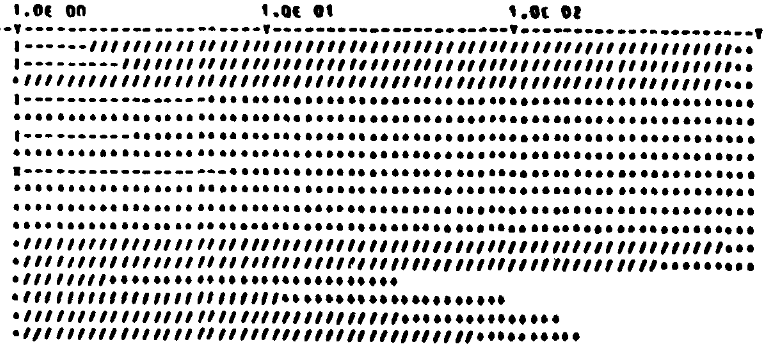

CURVE	MINIMUM	MAXIMUM
MIN	9.5278E-02	4.8195E 02
AVG	2.3556E-01	1.8474E 03
MAX	3.9955E-01	3.0825E 03

Figure 14-18. (Part 2 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

TOTAL TIME IN SUBDISPATCH QUEUE PLOT- 4

-= MIN /= AVG *= MAX

USERS	TIME	QUEUE STATISTICS #SD	CPU	TROUTS	I/O-BSY	LINE
2	238:19.56	242	16.19%	0.00%	0.00%	<TL>
1	239:19.56	44	18.23%	0.00%	34.09%	<TL>
1	240:19.56	30	14.89%	0.00%	70.00%	<TL>
2	241:19.56	133	42.67%	0.00%	57.89%	<TM>
2	242:19.56	2074	88.44%	97.44%	2.31%	<TM>
3	243:19.56	2582	88.23%	97.10%	0.89%	<TM>
4	244:19.56	3874	85.34%	97.83%	1.19%	<TM>
4	245:19.56	4809	85.14%	99.88%	0.00%	<TM>
5	246:19.56	4892	81.68%	94.71%	3.70%	<TO>
5	247:19.56	4692	79.00%	94.59%	1.47%	<TO>
5	248:19.56	4652	79.30%	95.79%	1.14%	<TO>
5	249:19.56	4755	77.43%	92.77%	4.40%	<TO>
5	250:19.56	4826	82.77%	97.10%	0.54%	<TM>
5	251:19.56	4669	74.03%	88.41%	5.42%	<TO>
6	252:19.56	4803	84.35%	99.35%	0.02%	<TM>
6	253:19.56	4832	82.95%	98.97%	0.56%	<TM>
6	254:19.56	4970	59.36%	99.40%	0.46%	<TM>
5	255:19.56	4800	59.35%	98.85%	0.40%	<TO>
5	256:19.56	4834	57.61%	99.81%	0.08%	<TN>
5	257:19.56	4850	59.71%	99.81%	0.08%	<TN>
5	258:19.56	4800	59.53%	99.65%	0.15%	<TO>
5	259:19.56	4908	59.46%	99.67%	0.20%	<TM>
6	300:19.56	5036	50.78%	98.29%	0.42%	<TM>
6	301:19.56	5096	36.19%	100.00%	0.00%	<TO>
7	302:19.56	5088	36.00%	99.49%	0.02%	<TQ>
7	303:19.56	4876	37.28%	99.51%	0.08%	<TN>
7	304:19.56	5131	31.62%	99.30%	0.23%	<TO>
9	305:19.56	5133	30.58%	98.97%	0.08%	<TN>
9	306:19.56	5284	30.00%	95.29%	3.69%	<T2>
9	307:19.56	4817	29.56%	96.74%	0.00%	<TP>
9	308:19.56	5136	27.79%	92.17%	3.54%	<T2>
8	309:19.56	4848	25.94%	90.00%	0.00%	<TM>
9	310:19.56	4865	29.75%	95.11%	0.16%	<TO>
9	311:19.56	5277	27.42%	95.71%	3.83%	<TM>
9	312:19.56	5032	24.31%	93.90%	0.00%	<TP>
9	313:19.56	5149	23.59%	90.13%	2.72%	<TP>
9	314:19.56	5259	26.95%	98.61%	1.12%	<TO>
10	315:19.56	4998	27.88%	98.84%	0.26%	<TD>
10	316:19.56	5051	30.78%	96.22%	2.57%	<TM>
10	317:19.56	4932	28.48%	96.90%	1.78%	<TM>
10	318:19.56	4790	31.74%	95.01%	2.99%	<TO>
10	319:19.56	4697	25.53%	95.55%	0.36%	<TO>
10	320:19.56	4860	29.58%	91.81%	4.07%	<T2>

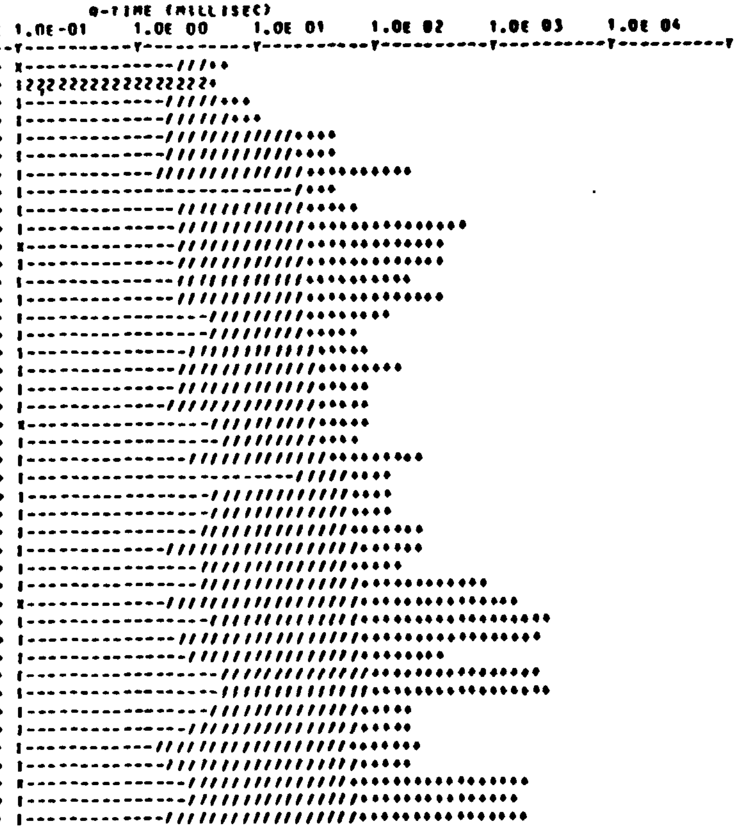

Figure 14-19. Total Time In Subdispatch Queue Report (Part 1 of 2)

USERS	TIME	RSD	CPU	TRDUTS	I/O-BSY	LINE
10	321:19.56	4792	27.88%	91.13%	0.48%	<T2>
10	322:19.56	4977	30.44%	93.09%	1.33%	<TM>
10	323:19.56	4452	29.71%	95.77%	3.35%	<T0>
10	324:19.56	5045	24.88%	98.68%	0.08%	<T0>
10	325:19.56	5017	28.70%	96.39%	2.25%	<T0>
10	326:19.56	5237	25.56%	96.79%	1.83%	<TN>
10	327:19.56	5137	29.44%	93.07%	1.54%	<TP>
10	328:19.56	4791	25.26%	93.54%	0.48%	<T0>
10	329:19.56	5317	26.28%	96.52%	3.23%	<TP>
10	330:19.56	5274	26.44%	98.24%	1.54%	<T0>
9	331:19.56	5389	26.03%	96.29%	3.58%	<TN>
10	332:19.56	5055	34.40%	96.85%	0.51%	<T0>
10	333:19.56	3524	66.87%	92.03%	4.63%	<T2>
10	334:19.56	2244	83.71%	86.27%	12.97%	<T0>
10	335:19.56	4834	48.86%	97.48%	1.99%	<T0>
10	336:19.56	5114	32.52%	92.47%	5.96%	<T0>
10	336:56.37	2982	38.67%	98.29%	0.54%	<TN>

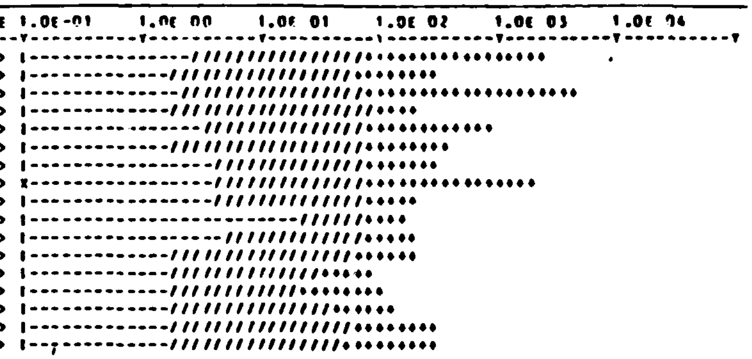

CURVE	MINIMUM	MAXIMUM
MIN	1.4010E 00	2.1625E 01
AVG	3.9045E 00	8.1607E 01
MAX	4.4690E 00	5.2880E 03

Figure 14-19. (Part 2 of 2)

This first SUBDISPATCH report is the sum of the next two reports. As with the RESPONSE TIME reports just described, the three are a set, and are meant to be evaluated together. Unfortunately, this is difficult for the same reasons.

14.6.7.4.2.6 Time in Subdispatch Queue Waiting Service. This report is shown in figure 14-20, and provides a measure of the amount of time spent by a subsystem (Sub-dispatch entry) in the Fault queue waiting for TSS action. Note that the anomaly discovered in section 14.6.7.4.2.2 does not appear here. This would indicate that TSS executive processing is not a factor in the delay incurred.

Comparisons of the time spent waiting TSS executive action cannot be made against any other measure or report currently reported by TEARS, but should be made across time. If, for example, a time period is discovered during which the average time a subsystem had to wait for executive service increased, analysts could search for some factor which is interfering with processor dispatches to the executive. Finally, this report should be used to discover whether at any time during the tracing interval, the total wait time exceeded the length of a TSS priority 'B' dispatch.

The overall wait time for servicing a Fault queue entry can be affected by increasing the priority 'B' parameters of TSS, as described in section 14.6.7.3.

14.6.7.4.2.7 Processor Time In Subdispatch. This report is shown in figure 14-21, and is the second subset of the TOTAL report described in section 14.6.7.4.2.5. The time interval shown is that spent by the subsystem from the time it was placed in the Sub-dispatch Ready queue, until it was retrieved by TSS from the Fault queue.

In the example shown, note the static nature of this average. Apparently, little else was going on in the system at the time. The average time spent by an entry in the Ready queue did not vary with the number of users in execution, as would have been the case if .MDISP had outstanding batch work.

14.6.7.4.2.8 CPU Monitor Driven Reporting. These reports, which are produced during the response phase of the TEARS System, have not been evaluated in this report, as their generation is triggered by trace records generated by the GMF/CPUM, not the TSSM.

14.6.7.4.2.9 TSS Subtraces Encountered Report. The usefulness of this report, shown in figure 14-22 is in the types of events which are reported. Along with the first RESPONSE phase report, this can be a useful indication of problems which might not be even implied by earlier reports.

The following events signal some problem which is affecting response time, and which should not be present in a healthy system.

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

TIME IN SUBDISPATCH QUEUE WAITING SERVICE PLOT- 3

* MIN /= AVG += MAX

USERS	TIME	QUEUE STATISTICS #SD	CPU	TROUTS	I/O-BSY	LINE
2	238:19.56	242	16.19%	0.00%	0.00%	<TL>
1	239:19.56	44	18.23%	0.00%	34.09%	<TL>
1	240:19.56	30	34.89%	0.00%	70.00%	<TL>
2	241:19.56	133	42.67%	0.00%	57.89%	<TM>
2	242:19.56	2074	88.44%	97.44%	2.31%	<TL>
3	243:19.56	2582	88.23%	97.10%	0.89%	<TM>
4	244:19.56	3974	85.34%	97.83%	1.19%	<TM>
4	245:19.56	4809	85.14%	99.88%	0.00%	<TO>
5	246:19.56	4892	81.68%	94.71%	3.70%	<TP>
5	247:19.56	4692	79.00%	94.59%	1.47%	<TO>
5	248:19.56	4652	79.30%	95.79%	1.14%	<TM>
5	249:19.56	4755	77.43%	92.77%	4.40%	<TO>
5	250:19.56	4826	82.77%	97.10%	0.54%	<TM>
5	251:19.56	4669	74.03%	88.41%	5.42%	<TO>
6	252:19.56	4803	84.55%	99.35%	0.02%	<TO>
6	253:19.56	4832	82.95%	98.97%	0.56%	<TM>
6	254:19.56	4970	59.36%	99.40%	0.46%	<TM>
5	255:19.56	4800	59.35%	98.85%	0.40%	<TO>
5	256:19.56	4834	59.61%	99.81%	0.08%	<TN>
5	257:19.56	4950	57.71%	97.81%	0.08%	<TN>
5	258:19.56	4800	57.53%	99.65%	0.15%	<TN>
5	259:19.56	4908	57.46%	99.67%	0.20%	<TM>
6	300:19.56	5036	50.78%	98.29%	0.42%	<TM>
6	301:19.56	5076	36.19%	100.00%	0.00%	<TO>
7	302:19.56	5088	36.00%	99.49%	0.02%	<TM>
7	303:19.56	4876	37.28%	99.51%	0.08%	<TM>
7	304:19.56	5131	31.62%	99.10%	0.23%	<TO>
8	305:19.56	5133	30.58%	98.97%	0.08%	<TN>
9	306:19.56	5284	30.00%	95.29%	3.69%	<T2>
9	307:19.56	4917	29.56%	96.74%	0.00%	<TP>
9	308:19.56	5136	27.79%	92.17%	3.54%	<T2>
8	309:19.56	4848	25.94%	90.00%	0.00%	<TN>
9	310:19.56	4865	29.75%	95.31%	0.16%	<TO>
9	311:19.56	5297	27.42%	95.71%	3.83%	<T3>
9	312:19.56	5037	24.31%	93.90%	0.00%	<TM>
9	313:19.56	5148	23.59%	90.13%	2.72%	<TP>
9	314:19.56	5259	26.95%	98.61%	1.12%	<TO>
10	315:19.56	4998	27.88%	98.86%	0.26%	<TO>
10	316:19.56	5051	30.78%	96.22%	2.57%	<TM>
10	317:19.56	4737	28.48%	96.90%	1.78%	<TN>
10	318:19.56	4790	31.74%	95.01%	2.99%	<TO>
10	319:19.56	4677	25.53%	95.55%	0.36%	<TO>
10	320:19.56	4860	29.58%	91.81%	4.07%	<TO>

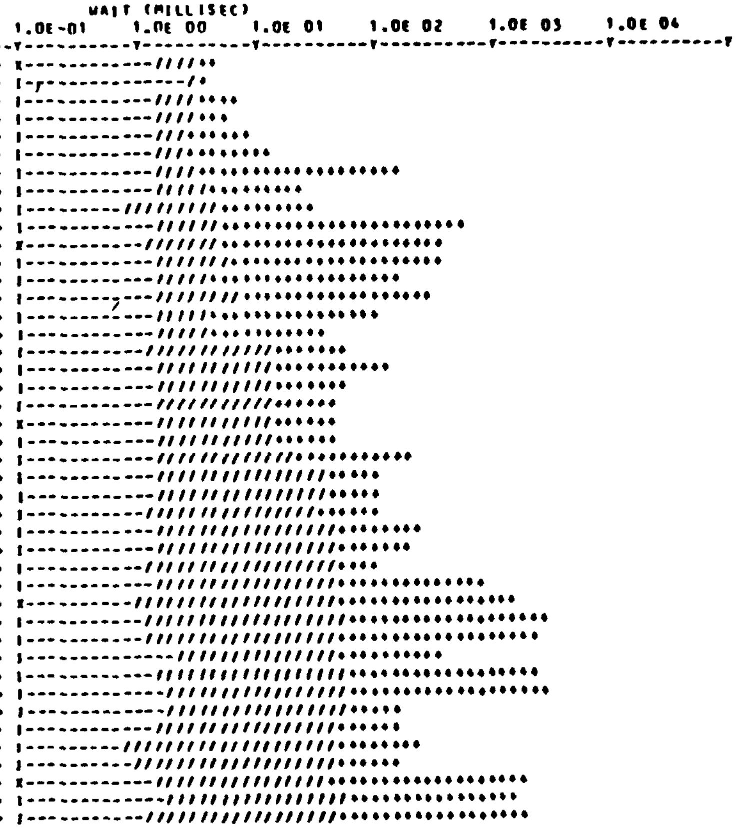

Figure 14-20. **Time In Subdispatch Queue Waiting Service Report (Part 1 of 2)**

USERS	TIME	BSD	CPU	TROUTS	I/O-BSY	LINE
10	321:19.56	4792	27.88%	91.13%	0.48%	<T2>
10	322:19.56	4977	30.44%	93.09%	1.33%	<TB>
10	323:19.56	4652	29.71%	95.77%	3.35%	<TQ>
10	324:19.56	5065	24.88%	98.68%	0.08%	<TO>
10	325:19.56	5017	28.70%	96.39%	2.25%	<TO>
10	326:19.56	5237	25.56%	96.79%	1.83%	<TN>
10	327:19.56	5137	29.44%	93.07%	1.54%	<TB>
10	328:19.56	4981	25.26%	93.54%	0.48%	<TQ>
10	329:19.56	5317	26.28%	96.52%	3.23%	<TP>
10	330:19.56	5276	26.44%	98.24%	1.54%	<TB>
9	331:19.56	5389	26.03%	96.27%	3.58%	<TN>
10	332:19.56	5055	34.40%	96.85%	0.51%	<TB>
10	333:19.56	3524	66.87%	92.03%	4.63%	<TN>
10	334:19.56	2244	83.71%	86.27%	12.97%	<TQ>
10	335:19.56	4836	49.86%	97.48%	1.99%	<TQ>
10	336:19.56	5114	32.52%	92.47%	5.96%	<TQ>
10	336:56.37	2992	38.67%	98.29%	0.54%	<TN>

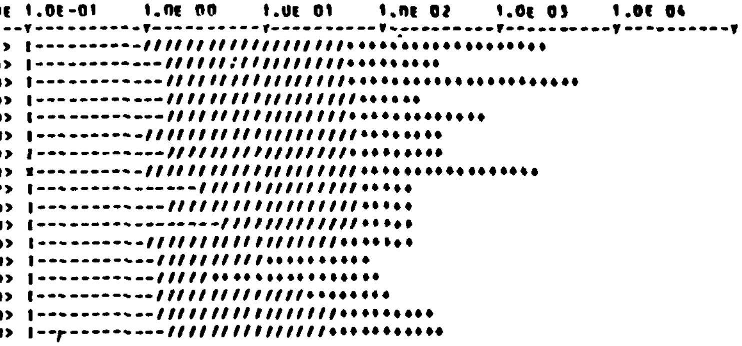

CURVE	MINIMUM	MAXIMUM
MIN	7.4500E-01	4.0370E 00
AVG	2.5872E 00	6.2100E 01
MAX	4.0930E 00	5.2661E 03

Figure 14-20. (Part 2 of 2)

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982
                         DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0237:19.557 FRI 83-02-11.  INITIAL GMC TAPE #0276


                            PROCESSOR TIME IN SUBDISPATCH                                                       PLOT- 6

                                                   -= MIN    /= AVG    += MAX

                          QUEUE STATISTICS                                       PROCTM (MILLISEC)
USERS    TIME             #SD    CPU       TROUTS  I/O-BSY  LINE 1.0E-01                   1.0E 00                1.0E 01
-------------------------------------------------------------------------Y-----------------------Y----------------------Y----------------------Y
  2   238:19.56           242   16.19%     0.00%    0.00%   <TL> X----------//////////++
  1   239:19.56            44   18.23%     0.00%   34.09%   <TL> I-T-------//////////////++++++++++
  1   240:19.56            30   34.89%     0.00%   70.00%   <TL> I----------/////////////////////+++++++
  2   241:19.56           133   42.67%     0.00%   57.89%   <TM> I--------/////////////////////////////+++++++++++
  2   242:19.56          2074   88.44%    97.44%    2.31%   <TM> I--------////////////////////////////////////////////////////+++++++
  1   243:19.56          2582   88.23%    97.10%    0.89%   <TM> I--------////////////////////////////////////////////////////+++++++
  4   244:19.56          3974   85.34%    97.83%    1.19%   <TM> I--------////////////////////////////////////////////////////+++++++
  4   245:19.56          4809   85.14%    99.88%    0.00%   <TM> I222222222222222222222222222222222222222222222222222222222222+++++++
  5   246:19.56          4892   81.68%    94.71%    3.70%   <TM> I--------////////////////////////////////////////////////////+++++++
  5   247:19.56          4692   79.00%    94.59%    1.47%   <TM> I--------////////////////////////////////////////////////////+++++++
  5   248:19.56          4652   79.30%    95.79%    1.14%   <TO> X--------////////////////////////////////////////////////////+++++++
  5   249:19.56          4755   77.43%    92.77%    4.40%   <TM> I-------/////////////////////////////////////////////////////+++++++
  5   250:19.56          4926   82.77%    97.10%    0.54%   <TM> I--------////////////////////////////////////////////////////+++++++
  5   251:19.56          4669   74.03%    88.41%    5.42%   <TO> I--------////////////////////////////////////////////////////+++++++
  6   252:19.56          4903   84.55%    99.35%    0.02%   <TM> I--------////////////////////////////////////////////////////+++++++
  6   253:19.56          4832   82.95%    98.97%    0.56%   <TO> I--------////////////////////////////////////////////////////+++++++
  6   254:19.56          4970   59.36%    99.40%    0.46%   <TM> I--------////////////////////////////////////////////////////+++++++
  5   255:19.56          4800   59.35%    98.85%    0.40%   <TO> I-------/////////////////////////////////////////////////////+++++++
  5   256:19.56          4834   59.61%    99.81%    0.08%   <TO> I----------//////////////////////////////////////////////////+++++++
  5   257:19.56          4850   59.71%    99.81%    0.08%   <TN> I----------//////////////////////////////////////////////////+++++++
  5   258:19.56          4800   59.53%    99.65%    0.15%   <TM> X----------//////////////////////////////////////////////////+++++++
  5   259:19.56          4909   59.46%    99.67%    0.20%   <TM> I--------////////////////////////////////////////////////////+++++++
  6   300:19.56          5036   50.78%    98.29%    0.42%   <TO> I-------/////////////////////////////////////////////////////+++++++
  6   301:19.56          5096   36.19%   100.00%    0.00%   <TM> I222222222222222222222222222222222222222222222222222222222222+++++++
  7   302:19.56          5088   36.00%    99.49%    0.02%   <TM> I--------////////////////////////////////////////////////////+++++++
  7   303:19.56          4876   37.28%    99.51%    0.09%   <TM> I--------////////////////////////////////////////////////////+++++++
  7   304:19.56          5131   31.62%    99.30%    0.23%   <TO> I--------////////////////////////////////////////////////////+++++++
  8   305:19.56          5133   30.59%    98.97%    0.08%   <TP> I-----///////////////////////////////////////////////////////+++++++
  9   306:19.56          5284   30.00%    95.29%    3.69%   <TP> I-----///////////////////////////////////////////////////////+++++++
  9   307:19.56          4817   29.56%    96.74%    0.00%   <TP> I-------/////////////////////////////////////////////////////+++++++
  9   308:19.56          5136   27.79%    92.17%    3.54%   <TQ> X-----///////////////////////////////////////////////////////+++++++
  8   309:19.56          4848   25.94%    90.00%    0.00%   <TN> I-----///////////////////////////////////////////////////////+++++++
  9   310:19.56          4865   29.75%    95.31%    0.16%   <TM> I-----///////////////////////////////////////////////////////+++++++
  9   311:19.56          5297   27.42%    95.71%    3.83%   <TM> I-----///////////////////////////////////////////////////////+++++++
  9   312:19.56          5032   24.31%    93.90%    0.00%   <TP> I--------////////////////////////////////////////////////////+++++++
  9   313:19.56          5148   23.59%    90.13%    2.72%   <TP> I-----///////////////////////////////////////////////////////+++++++
  9   314:19.56          5259   26.95%    98.61%    1.12%   <TN> I-----///////////////////////////////////////////////////////+++++++
 10   315:19.56          4998   27.88%    98.84%    0.26%   <TQ> I-----///////////////////////////////////////////////////////+++++++
 10   316:19.56          5051   30.78%    96.22%    2.57%   <TN> I-----///////////////////////////////////////////////////////+++++++
 10   317:19.56          4932   28.48%    96.90%    1.78%   <T2> I-----///////////////////////////////////////////////////////+++++++
 10   318:19.56          4790   31.74%    95.01%    2.99%   <TN> X-----///////////////////////////////////////////////////////+++++++
 10   319:19.56          4697   25.53%    95.55%    0.36%   <TO> I--------////////////////////////////////////////////////////+++++++
 10   320:19.56          4860   29.58%    91.81%    4.07%   <TN> I-----///////////////////////////////////////////////////////+++++++
```

Figure 14-21. **Processor Time In Subdispatch (Part 1 of 2)**

USERS	TIME	#SD	CPU	TROUTS	I/O-BSY	LINE
10	321:19.56	4792	27.88%	91.13%	0.48%	<TP>
10	322:19.56	4977	30.44%	93.09%	1.33%	<TO>
10	323:19.56	4852	29.71%	95.77%	3.35%	<TQ>
10	324:19.56	5065	24.88%	98.68%	0.06%	<TO>
10	325:19.56	5017	28.70%	96.39%	2.25%	<TN>
10	326:19.56	5237	25.56%	96.79%	1.83%	<T2>
10	327:19.56	5137	29.44%	93.07%	1.54%	<TQ>
10	328:19.56	4981	25.26%	93.54%	0.48%	<TN>
10	329:19.56	5317	26.28%	96.52%	3.23%	<TD>
10	330:19.56	5274	26.44%	98.24%	1.54%	<TM>
9	331:19.56	5389	26.03%	96.29%	3.58%	<TM>
10	332:19.56	5055	34.40%	96.85%	0.51%	<TN>
10	333:19.56	3526	66.87%	92.03%	4.63%	<TM>
10	334:19.56	2244	83.71%	86.27%	12.97%	<TD>
10	335:19.56	4834	48.86%	97.48%	1.99%	<TD>
10	336:19.56	5116	32.52%	92.47%	5.96%	<TO>
10	336:56.37	2982	38.67%	98.29%	0.54%	<TQ>

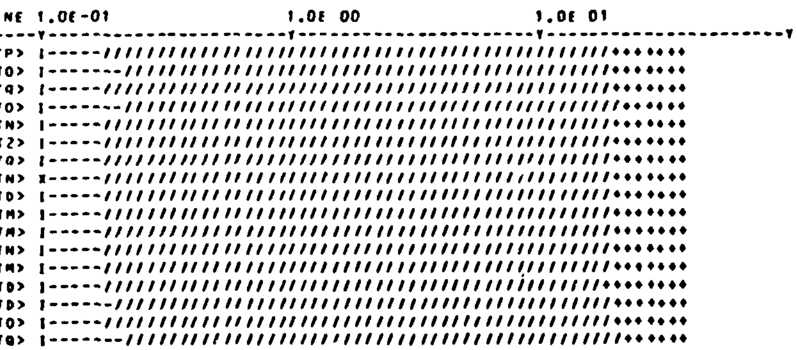

CURVE	MINIMUM	MAXIMUM
MIN	1.7100E-01	2.0187E 01
AVG	6.3212E-01	2.0356E 01
MAX	7.5000E-01	4.0296E 01

Figure 14-21. (Part 2 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM ASCC1 AT 0257:19.557 FRI 83-02-11, INITIAL GRC TAPE #0276

TSS SUBTRACES ENCOUNTERED (LOW HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	HISTOGRAM 1
12	12	0.001	0.001	1-	1	I	LOG-ON COMPLETE
0	12	0.001	0.	2-	2	I	SWAP SPACE DENIED
>44	56	0.005	0.004	3-	3	I	PAT DENIAL
44	100	0.007	0.004	4-	4	I	GETSYE DEALLOCATE
44	144	0.012	0.004	5-	5	I	DEALLOC COMPLETE
7	151	0.013	0.001	6-	6	I	ERROR MESSAGE
160	311	0.026	0.014	7-	7	I	POPUP PRIMITIVE
279	590	0.050	0.024	8-	8	I	CALLP PRIMITIVE
145	735	0.063	0.012	9-	9	I	ENTER BUILD MODE
190	925	0.079	0.016	10-	10	I	EXEC PRIMITIVE
31	956	0.081	0.003	11-	11	I	SYSTM PRIMITIVE
4	960	0.082	0.000	12-	12	I	LOG-OFF
110	1070	0.091	0.009	13-	13	I	COMMAND RECEIVED
6331	7401	0.630	0.539	14-	14	IX	PERIODIC CHECK
0	7401	0.630	0.	15-	15	I	GEWAKE NO USERS
8	7409	0.631	0.001	16-	16	I	BREAK/DISCONNECT
133134	140543	11.972	11.341	17-	17	IXXXXXXXXXXX	GEWAKE S/D DONE
13657	154200	13.135	1.163	18-	18	IX	GEWAKE S/D BUSY
5	154205	13.136	0.000	19-	19	I	APR RIO C/C
185	154390	13.151	0.016	20-	20	I	BUILD INPUT REC
0	154390	13.151	0.	21-	21	I	ST** I/O COMPLETE
3	154393	13.152	0.000	22-	22	I	RECONNECT MODE
10442	164835	14.041	0.889	23-	23	IX	PROCESS DRL
4061	168896	14.387	0.346	24-	24	I	REQUEST FILE I/O
3912	172808	14.720	0.333	25-	25	I	DISK I/O COMPLETE
1	172809	14.720	0.000	26-	26	I	DRL DEFIL DENIAL
248	173057	14.741	0.021	27-	27	I	ISSUE MME GEFSYE
44	173101	14.745	0.004	28-	28	I	BAD FILACT PARAM
248	173349	14.766	0.021	29-	29	I	GEFSYE/27 COMPLETE
0	173349	14.766	0.	30-	30	I	DELAY 200 MS
0	173349	14.766	0.	31-	31	I	DELAY 2 SEC
12	173361	14.767	0.001	32-	32	I	STORE USERID
0	173361	14.767	0.	33-	33	I	PASUST -1 COMP
1	173362	14.767	0.000	34-	34	I	DRL MORLNK ERROR
0	173362	14.767	0.	35-	35	I	SPAWN I/O COMPLETE
0	173362	14.767	0.	36-	36	I	BATCH JOB
0	173362	14.767	0.	37-	37	I	DRL SPAWN ERROR
0	173362	14.767	0.	38-	38	I	WRITE ABRT
0	173362	14.767	0.	39-	39	I	ABRT I/O COMPLETE
30	173392	14.770	0.003	40-	40	I	EXEC* - DRL RESTOR
30	173422	14.773	0.003	41-	41	I	LOAD DRL RESTOR
49	173471	14.777	0.004	42-	42	I	RESTOR LOAD COMPL
38	173509	14.780	0.003	43-	43	I	PERM I/O RESTOR
0	173509	14.780	0.	44-	44	I	START LINE SWITCH
1	173510	14.780	0.000	45-	45	I	FILE GROW .MFS19
1	173511	14.780	0.000	46-	46	I	GROW COMPLETE
0	173511	14.780	0.	47-	47	I	CONSOLE I/O
0	173511	14.780	0.	48-	48	I	DRL JOUT RETURN
0	173511	14.780	0.	49-	49	I	JOUT OUTPUT BUSY
31	173542	14.783	0.003	50-	50	I	START GWAKE

1173949 ENTRIES TOTAL   1003407 OUT OF RANGE

Figure 14-22. TSS Subtraces Encountered (Part 1 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1992

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

TSS SUBTRACES ENCOUNTERED (HIGH HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	HISTOGRAM 2
173542	173542	14.783	14.783	1- 54	IXXXXXXXXXXXXXX	DRL SAVE COMPLETE
0	173542	14.783	0.	55- 55	I	IDS ATTRI .MFS11
0	173542	14.783	0.	56- 56	I	IDS ATTRI STATUS
0	173542	14.783	0.	57- 57	I	DRL T.STOT DENIAL
0	173542	14.783	0.	58- 58	I	UST SWITCH T.CONN
0	173542	14.783	0.	59- 59	I	TAP* I/O COMPLETE
8855	182397	15.537	0.754	60- 60	IX	ALLOCATOR SERVICES
420242	602639	51.334	35.797	61- 61	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	ENTER PROC ALLOC
6190	608829	51.862	0.527	62- 62	IX	ENTER MEMORY ALLOC
1042	609871	51.950	0.089	63- 63	I	ENTER SWAP DECIS
1153	611024	52.049	0.098	64- 64	I	NEXT PRIMITIVE IS
30	611054	52.051	0.003	65- 65	I	CONSIDER TSS INCR
12	611066	52.052	0.001	66- 66	I	INITIATE SIZE INCR
93	611159	52.060	0.008	67- 67	I	SET UP FENCE
3	611162	52.060	0.000	68- 68	I	MEMORY RESERVE
784	611946	52.127	0.067	69- 69	I	FORCE SWAP
0	611946	52.127	0.	70- 70	I	TERMINATE SWAP
2	611948	52.127	0.000	71- 71	I	TSSL UST AREA +1K
59	612007	52.132	0.005	72- 72	I	GEMORE MEMORY
5	612012	52.133	0.000	73- 73	I	GEMORE SUCCESSFUL
3	612015	52.133	0.000	74- 74	I	MEMORY RELEASE
172	612187	52.148	0.015	75- 75	I	GEMORE REFUSED
3	612190	52.148	0.000	76- 76	I	CANCEL CRUN
2567	614757	52.367	0.219	77- 77	I	ATTEMPT MEM ALLOC
3072	617829	52.628	0.262	78- 78	I	MEM MAP CHANGE
1090	618919	52.721	0.093	79- 79	I	SWAP PROGRAM
5685	624604	53.205	0.484	80- 80	I	TIME TYPES
212	624816	53.223	0.018	81- 81	I	REQUEST EOM
141	624957	53.235	0.012	82- 82	I	SCAN COMMAND LIST
19	624976	53.237	0.002	83- 83	I	EOM REFUSED
180	625156	53.252	0.015	84- 84	I	INPUT COMPLETE
18	625174	53.254	0.002	85- 85	I	POPUP TO CALLSS
1	625175	53.254	0.000	86- 86	I	APB COMPLETE
919	626094	53.332	0.078	87- 87	I	TERM I/O COMPLETE
266036	892130	75.994	22.662	88- 88	IXXXXXXXXXXXXXXXXXXXXXX	S/D FAULT QUEUE
7425	899555	76.626	0.632	89- 89	IX	PROCESS .CRTSO
3	899558	76.627	0.000	90- 90	I	TRACE ON/OFF
266042	1165600	99.289	22.662	91- 91	IXXXXXXXXXXXXXXXXXXXXXX	MAKE S/D ENTRY
12	1165612	99.290	0.001	92- 92	I	LOG-ON REQUEST
0	1165612	99.290	0.	93- 93	I	BAD LINE STATUS
12	1165624	99.291	0.001	94- 94	I	CHECK VIP TERM TYP
0	1165624	99.291	0.	95- 95	I	UST WAIT TIME
0	1165624	99.291	0.	96- 96	I	REJECT USER
1	1165625	99.291	0.000	97- 97	I	UST COMPRESSION
0	1165625	99.291	0.	98- 98	I	TSSN UST AREA +1K
0	1165625	99.291	0.	99- 99	I	UST AREA -1K
6680	1172305	99.860	0.569	100- 100	IX	TERMINAL I/O
0	1172305	99.860	0.	101- 101	I	CRUN *** FUNC
12	1172317	99.861	0.001	102- 102	I	REMOTE INQUIRY
3	1172320	99.861	0.000	103- 103	I	VIP INPUT COMPLETE
1629	1173949	100.000	0.139	104- 104	I	DRL COMPLETE

1173949 ENTRIES TOTAL 0 OUT OF RANGE

Figure 14-22. (Part 2 of 2)

Trace type 3. PAT denial. This is caused by a lack of PAT space in TSS slave service areas. The occurrence of this type of trace means that too few SSAs are being reserved by TSS because of either disk fragmentation, heavy user load, or some anomaly in the number of files used by TSS terminals. In any case, the solution lies in raising the number of SSAs which are reserved by TSS. Section 14.6.7.3 describes a patch to location 136 of TSSA which should be used to prevent this trace from occurring.

Trace type 31. This trace is generated when a user subsystem is delayed two seconds due to SMC contention problems. The cataloged permanent disc files are stored under USERIDS which are divided into SMC sections. When one user accesses an SMC section, the software will momentarily lock it. If this trace type occurs more than a few times, the user affected will begin to note a delay in his processing. In these cases, analysis is indicated to find how this contention can be minimized or prevented.

Trace type 93. This trace type should not occur. If it is discovered, a hardware or GRTS problem is indicated which should be reported to field engineering personnel.

Trace type 83. This trace type indicates an overload in the subsystem memory area which is delaying VIP terminal output from being produced. If discovered, either the minimum size for TSS should be modified upwards, or the maximum number of VIP terminals allowed should be decreased.

Trace types 65, 66, 67, 68, 69, 70, 72, and 73. High values for these trace types indicate that the subsystem memory area is facing very high contention, causing urgent (overdue) memory allocations, force swaps, and TSS size increases. By raising the maximum (and minimum) TSS sizes, and by lowering the 'large' subsystem definition, the occurrence of these trace types should be minimized.

In the example shown, note that a PAT denial occurred 44 times during the tracing interval. An EBM buffer for VIP output was refused 19 times. TSS considered a size increase 59 times, (trace type 65), and initiated such an increase 12 times (trace type 66). 93 urgent allocations occurred, (trace type 67, 68), and 784 force swaps were made (trace type 69).

14.6.7.4.2.10 Derails Report. This final report produced during the RESPONSE phase of TEARS processing, shown in figure 14-23, is a statistical analysis of the number of times each Derail type is executed. It is of an academic interest only, and provides little information useful for analysis of response time problems.

14.6.7.4.3 Analysis of TEARS Emulation Phase Reports. Once the TEARS user has analysed the reports produced during the RESPONSE phase of reduction, and identified the periods in which he has interest, he should produce the EMULATION reports for these periods. The fields of interest within the reports produced during this second phase of data reduction are described in this section.

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #D276

DERAILS (LOW HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS DRL	#	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	HISTOGRAM 3
						I....I....I....I....I....I....I....I....I....I....I	
36	36	0.345	0.345	-3-	-3	I	W.STOR
0	36	0.345	0.	-2-	-2	I	W.ARCH
37	73	0.699	0.354	-1-	-1	I	W.USER
0	73	0.699	0.	0-	0	I	ZERO
3914	3987	38.182	37.483	1-	1	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	DIO
3990	7977	76.393	38.211	2-	2	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	KOUT
179	8156	78.105	1.714	3-	3	IXX	KOUTN
286	8442	80.847	2.739	4-	4	IXXX	KIN
151	8593	82.293	1.446	5-	5	IX	RETURN
180	8773	84.016	1.724	6-	6	IXX	DEFIL
0	8773	84.016	0.	7-	7	I	ABORT
187	8960	85.807	1.791	8-	8	IXX	SETSWN
38	8998	86.171	0.364	9-	9	I	RSTSWN
230	9228	88.374	2.203	10-	10	IXX	REW
2	9230	88.393	0.019	11-	11	I	FILSP
92	9322	89.274	0.881	12-	12	IX	RETFIL
35	9357	89.609	0.335	13-	13	I	RELMEM
22	9379	89.820	0.211	14-	14	I	ADDMEM
4	9383	89.859	0.038	15-	15	I	CONFIL
1	9384	89.868	0.010	16-	16	I	SNUMB
35	9419	90.203	0.335	17-	17	I	TIME
4	9423	90.241	0.038	18-	18	I	PASAFT
40	9463	90.624	0.383	19-	19	I	TERMTP
0	9463	90.624	0.	20-	20	I	PDIO
49	9512	91.094	0.469	21-	21	I	RESTOR
0	9512	91.094	0.	22-	22	I	SPAWN
0	9512	91.094	0.	23-	23	I	TAPEIN
23	9535	91.314	0.220	24-	24	I	CALLSS
12	9547	91.429	0.115	25-	25	I	USERID
4	9551	91.467	0.038	26-	26	I	TERMPG
145	9696	92.856	1.389	27-	27	IX	PASUST

10442 ENTRIES TOTAL 746 OUT OF RANGE

Figure 14-23. **Derails Report (Part 1 of 2)**

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0237:19.557 FRI 83-02-11. INITIAL GMC TAPE #0276

DETAILS (HIGH HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS BRL	#	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	HISTOGRAM &
						I....I....I....I....I....I....I....I....I....I....I	
9696	9696	92.856	92.856	-3-	28	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	MORLNK
0	9696	92.856	0.	29-	29	I	NEWUSR
292	9988	95.652	2.796	30-	30	IX	FILACT
0	9988	95.652	0.	31-	31	I	SETLNO
0	9988	95.652	0.	32-	32	I	SYSRET
0	9988	95.652	0.	33-	33	I	STPSYS
0	9988	95.652	0.	34-	34	I	STATUS
0	9988	95.652	0.	35-	35	I	BRLDSC
2	9990	95.671	0.019	36-	36	I	PASDES
51	10041	96.160	0.488	37-	37	I	JSTS
12	10053	96.275	0.115	38-	38	I	CGROUT
4	10057	96.313	0.038	39-	39	I	PART
13	10070	96.437	0.124	40-	40	I	GROW
0	10070	96.437	0.	41-	41	I	ABTJOB
0	10070	96.437	0.	42-	42	I	CONSOL
49	10119	96.907	0.469	43-	43	I	SWITCH
15	10134	97.050	0.144	44-	44	I	BRLIMT
0	10134	97.050	0.	45-	45	I	JOUT
21	10155	97.251	0.201	46-	46	I	KOTNOW
49	10204	97.721	0.469	47-	47	I	OBJTIM
0	10204	97.721	0.	48-	48	I	PASFLR
0	10204	97.721	0.	49-	49	I	STOPPT
0	10204	97.721	0.	50-	50	I	BRLSAV
0	10204	97.721	0.	51-	51	I	TASK
75	10279	98.439	0.718	52-	52	I	PSEUDO
12	10291	98.554	0.115	53-	53	I	PRGDES
33	10324	98.870	0.316	54-	54	I	GWAKE
0	10324	98.870	0.	55-	55	I	T.CRYP
0	10324	98.870	0.	56-	56	I	ATTRI
0	10324	98.870	0.	57-	57	I	T.STAT
36	10360	99.215	0.345	58-	58	I	T.GOTO
38	10398	99.579	0.364	59-	59	I	T.CMOV
0	10398	99.579	0.	60-	60	I	T.LINL
0	10398	99.579	0.	61-	61	I	T.SYOT
0	10398	99.579	0.	62-	62	I	T.CONN
27	10425	99.837	0.259	63-	63	I	T.CFIO
0	10425	99.837	0.	64-	64	I	T.EXEC
0	10425	99.837	0.	65-	65	I	T.RSCC
0	10425	99.837	0.	66-	66	I	T.MAIL
17	10442	100.000	0.163	67-	67	I	T.ERR

10442 ENTRIES TOTAL   0 OUT OF RANGE

Figure 14-23. **(Part 2 of 2)**

14.6.7.4.3.1 Exception Message Report. The first report produced during the EMULATION phase is an index similar to the event log described earlier. Figure 14-24 illustrates how this report can display the problem discovered through examination of the RESPONSE phase reports in this section. Note that terminal TM does have a heavily processor bound subsystem in memory, but no other exceptions occurred. Several other types of exception can be displayed by this report, and each is described below.

SWAP SPACE DENIED. The TSS swap files are full, or not enough room remains on any one of them to hold the subsystem for this user. This may be due to the size of the subsystem being swapped (or, more frequently, pushed down). The inability to swap this user subsystem is not a denial to this user, who will continue to enjoy processor attention if required, but will affect other memory allocation requests. In any case, the resolution involves raising the maximum size of the TSS swap files, as described in section 14.6.7.3.

JOUT OUTPUT BUSY. Self explanatory.

CANNOT SWAP. Insufficient room on the device containing one of the swap files for TSS remains to allow that file to grow. Resolution involves somehow depopulating that device, or allowing the length of disk extents to be raised by recovering disk fragments by performing a cold boot.

I/O ERROR ON DEVICE. Self Explanatory. Resolution involves reformatting the device during a cold boot.

UNDEFINED DEVICE. This occurs when users attempt to restore files from systems having devices which are not present. Privileged FMS directives can be used to delete the file reference.

SEEK ERROR. Self Explanatory. This type of problem may indicate software, hardware, or firmware errors.

FAILURE IN NAME SCAN. Program stack errors.

BAD SPACE INVENTORY. Device error.

BAD SPACE TABLE. Device error.

UNACCOUNTABLE ERROR 1.

CHECKSUM ERROR. Device or I/O error.

PAGE REQUEST DEADLOCK. IDS error.

UNACCOUNTABLE ERROR 2.

BAD PRIMITIVE. GCOS problem, indicates some master mode code destroyed.

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0236:21.243 FRI 83-02-11. INITIAL GMC TAPE #0276

EXCEPTION MESSAGE REPORT

TOD & RECNO		EXCEPTION MESSAGE			
0241:32.735	5957	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	65219	
0241:33.749	6222	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	120427	....
0241:34.770	6490	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	176043	....
0241:35.782	6754	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	231577	....
0241:36.794	7020	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	287113	....
0241:37.809	7286	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	342653	....
0241:38.829	7552	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	398487	....
0241:39.840	7816	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	454023	....
0241:40.842	8081	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	508848	....
0241:41.854	8347	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	564388	....
0241:42.865	8608	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	619867	....
0241:43.869	8897	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	674343	....
0241:44.876	9305	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	728046	....
0241:45.895	9569	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	783881	....
0241:46.901	9840	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	839179	....
0241:47.913	10106	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	894714	....
0241:48.928	10370	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	950249	....
0241:49.946	10638	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1005898	....
0241:50.958	10902	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1061434	....
0241:51.970	11166	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1116969	....
0241:52.983	11436	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1172509	....
0241:54.000	11702	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1228123	....
0241:55.013	11968	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1283458	....
0241:56.024	12233	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1338873	....
0241:57.035	12497	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1394410	....
0241:58.054	12763	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1450037	....
0241:59.070	13033	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1505632	....
0242:00.082	13297	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1561170	....
0242:01.095	13563	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1616710	....
0242:02.113	13826	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1672175	....
0242:03.125	14090	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1727709	....
0242:04.138	14356	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1783244	....
0242:05.154	14626	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1838802	....
0242:06.170	14890	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1894338	....
0242:07.178	15149	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	1949069	....
0242:08.191	15417	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2004608	....
0242:09.202	15678	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2059729	....
0242:10.218	15944	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2115271	....
0242:11.220	16207	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2169933	....
0242:12.227	16460	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2224504	....
0242:13.249	16710	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2280309	....
0242:14.272	16968	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2336035	....
0242:15.289	17222	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2391565	....
0242:16.309	17474	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2447094	....
0242:17.313	17721	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2501903	....
0242:18.330	17973	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2557433	....
0242:19.348	18225	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2612968	....
0242:20.364	18479	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2668499	....
0242:21.383	18726	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2723592	....
0242:22.398	18978	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2779121	....
0242:23.403	19233	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2833462	....
0242:24.419	19483	LINE <TM>, USERID B2910PX0	!! CPU HOG? CPS SINCE SWAPIN =	2888990	....

Figure 14-24. Exception Message Report

BAD PROGRAM DESCRIPTOR. GCOS problem.

LOOP IN PRIMITIVES. GCOS problem.

PROGRAM TOO LARGE TO SWAP. The subsystem memory area assigned to this user is too large to fit on any of the TSS swap files. To resolve, increase the minimum and maximum sizes for these files.

BAD STATUS-SWAP OUT. I/O error.

BAD STATUS - SWAP IN. I/O error.

BAD STATUS - LOAD. I/O error on TSS program file containing subsystem being loaded.

ERROR IN WRITING TAP*. During TAPE mode processing for the user.

NOT ENOUGH CORE TO RUN PROGRAM. Produced after a settable interval has passed to prevent memory allocations from becoming urgent several times.

OUT OF SWAP SPACE. TSS has grown all of its swap files to the currently set maximum size.

SY** I/O ERROR I/O error on user current file.

DRL DEFIL - NO FILE SPACE. Self explanatory.

DRL DEFIL - NO PAT SPACE. TSS has filled its SSA space, and cannot define another file. Raise the number of SSAs to retain during startup.

DRL MORLNK - NO FILE SPACE. Self explanatory.

DRL MORLNK - PAT FULL. SSA space full.

DRL JOUT - BATCH SYSTEM FULL. No entries are available in GEOT queue. Transient problem not related to TSS.

DRL JOUT - LOST PAT. GCOS error.

DRL SPAWN - NO PAT. TSS SSA space full.

DRL SPAWN - SCHEDULER QUEUE FULL. Self explanatory. Transient problem.

DRL TASK - NO PROGRAM NO. All program numbers are in use in the system.

DRL TASK - BAD *J STATUS. I/O or device error.

DRL T.SYOT - SYSTEM LOADED. GEOT queue full. Transient error.

DRL T.SYOT - BACKDOOR FILE NOT CONFIGD. Subsystem software error.

GEMORE DENIED. Insufficient core memory for growth or impossible to release memory at this time. Transient problem.

CPU HOG? CPS SINCE SWAPIN = nnnnnn. This exception is produced when a subsystem seems to be processor bound. In the example shown in figure 14-24, terminal TM is executing a subsystem which is in this state. The periods to the right of this exception message are repetition factors which are used in place of displaying this message more than once each second.

14.6.7.4.3.2 USERS Map Report. The reader is referred to the GMF Users manual for an explanation of this report, which cannot be used to analyse or identify response time problems.

14.6.7.4.3.3 TSS Core Map Report. The reader is referred to the GMF Users Manual for an explanation of this report, which is also not useful in the identification or analysis or response time problems.

14.6.7.4.3.4 ERROR MESSAGE Report. This report is produced when errors in program logic prevent analysis of trace data collected by the GMF/TSSM. The EMULATION phase of TEARS is a primitive model of the ways in which TSS itself works. When the actual operation of the system differs from the model (and if this difference is noticed by TEARS) an error message is produced.

This report, as with the previous two reports, is not useful to analysts interested in the identification or analysis of response time problems.

14.6.7.4.3.5 INTERTRACE Reports. The purpose of these reports is to somehow track periods when the intervals between the occurrences of GMF/TSSM traces (i.e., their generation) seems to be large. A large (or long) interval between the occurrence of GMF/TSSM traces would imply that something is preventing the execution of the TSS executive. If the times when these intervals are large could be trapped and reported, some investigation could be performed to identify why the executive was not being dispatched to.

The implementation of this concept has been made as a function of the time delays between GMF/TSSM traces (as determined by the time stamp included with each trace record), with corrections made for periods when TSS relinquishes the processor voluntarily (GEWAKE). The concept of this type of reporting is valid, but cannot be collected from the TSS executive. This report is being examined to improve its usability.

14.6.7.4.3.6 Memory Activity Report. This report can be used to indicate the frequency and times during which activities occurred which are related to subsystem memory management. By documentation, this report should be produced each time an exception in processing is noted (section 14.6.7.4.3.1). Examination of the reports produced by the RESPONSE phase again indicate discrepancies with this report. These descrepancies are being investigated.

14.6.7.4.3.7 SUBDISPATCH Reports. This report is currently in error.

14.6.7.4.3.8 Users Swaps (Swap Rate) Report. This report is not useful in the identification or analysis of response time problems, and the user is referred to the GMF Users Manual for an explanation of its purpose.

14.6.7.4.3.9 User Swaps (Swap Amount). An example of this report is shown in figure 14-25. This is the best report currently available which can be used to identify the sizes of the subsystems which have become candidates for swap.

14.6.7.4.3.10 User Swaps (Duration) Report. The User Swaps (Duration) report describes the length of subsystem swaps, and is illustrated in figure 14-26. The user should examine the content of this report to gain an understanding of the volume of allocation work being faced by the system, as this is the factor which will influence the duration swaps for subsystems.

14.6.7.4.3.11 User I/Os - Duration Report. This report provides a measure of the time spent by users in a relinquished state pending disk I/O operations. There are no Tunable parameters associated with I/O operations which can affect response times. In discovering the source of a delay for a particular user, this report might become useful in the rare situation where an I/O operation was delayed for some reason. Figure 14-27 contains an example of this report.

The only conceivable way in which an I/O operation might be delayed for a significant period is if the hardware itself were at fault. If this were the case, alert messages would be placed on the system console.

14.6.7.4.3.12 In SMC Wait (Duration). As described earlier, the occurrence of the GMF/TSSM traces associated with SMC waiting are symptomatic of a contention between some TSS user and another process using the same SMC section. Figure 14-28 contains an example of this report, which can be useful to identify the extent of the delay incurred. Resolution of the problem indicated by these traces will involve some modification to the SMC section assignment of the users affected.

The trace information collected by the GMF/TSSM will be of only marginal value in the identification of the processes which are in contention, unless both are TSS users, which is not necessarily the case.

14.6.7.4.4 Resolution of Response Time Problems. The sections above have described the indications of response time problems which can be discovered through the review of the reports produced by the TEARS system. In general, the anomalies to the normally encountered values or curves are the items which should prompt attention or focus. If the problem is a continuing one,

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982
                         DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0236:21.243 FRI 83-02-11.  INITIAL GMC TAPE #0276


                                                   USER SWAPS (SWAP AMOUNT)                                              PLOT-12
                                                        -= AVG    /= TOT

                           CORE IN INTERVAL (K)                         SWAP SIZE (K)
 USERS   TIME          <2   <4   <8  <16  <32  <64  >64  LINE 1.0E 00           1.0E 01           1.0E 02           1.0E 03
------------------------------------------------------------------Y-----------------Y-----------------Y-----------------Y-----------------Y
   0  0237:19.55  DATA NOT COLLECTED IN THIS INTERVAL.           X
   2  0237:39.55   0    8                               <TL> I-----///////////////
   2  0237:59.55   0    7                               <TL> I-----//////////////
   2  0238:19.55   0    7                               <TL> I-----//////////////
   2  0238:39.55   2    3    1                          <TL> I-----/////////////
   2  0238:59.55                                             I
   1  0239:19.55   3    5                               <TL> I---///////////////
   1  0239:39.55   0    0    0    0    0    4           <TL> I---------------------------//////////
   1  0239:59.55                                             I
   1  0240:19.55                                             I
   2  0240:39.55   2    5    1                          <TM> X------///////////////
   2  0240:59.55                                             I
   2  0241:19.55   3    0    2    1                     <TM> I----------//////////////
   2  0241:39.55   1    0    1    1                     <TM> I---------------////////
   2  0241:59.55   0    0    0    0    0    2           <TL> I---------------------------/////
   2  0242:19.55                                             I
   3  0242:39.55   1    8    0    0    0    1           <TM> I------------/////////////////
   3  0242:59.55   0    0    0    0    0    2           <TM> I---------------------------/////
   3  0243:19.55                                             I
   4  0243:39.55   2    5    1                          <TO> I------///////////////
   4  0243:59.55   1    0    1    1    0    2           <TL> X--------------------////////////
   4  0244:04.13   0    0    0    0    0    1           <TM> I2222222222222222222222222


  CURVE       MINIMUM        MAXIMUM
  AVG       1.6250E 00     3.4250E 01
  TOT       1.3000E 01     1.3700E 02
```

Figure 14-25. User Swaps (Swap Amount) Report

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982
                    DISTRIBUTION COLLECTED ON SYSTEM BSCCT  AT 0236:21.243 FRI 83-02-11.  INITIAL GMC TAPE #0276


                                                                    USER SWAPS (DURATION)                                      PLOT-13

                                                                        -= MIN    /= AVG    *= MAX

                               DURATION IN INTERVAL (MS)                        SWAP TIMES (HUNDREDTHS SEC)
USERS   TIME          <1    <2    <4    <8   <16   <32   >32  LINE 1.0E-01            1.0E 00          1.0E 01          1.0E 02
----------------------------------------------------------------------Y-----------------Y-----------------Y-----------------Y----------------Y
  0   0237:19.55  DATA NOT COLLECTED IN THIS INTERVAL.                 I
  2   0237:39.55      0     0     3     1                        <TL>  I-----------------------///*****
  2   0237:59.55      0     0     3     1                        <TL>  I-----------------------//****
  2   0238:19.55      0     0     2     0     0     1            <TL>  I-----------------------//////////******
  2   0238:39.55      0     0     3     2                        <TL>  I---------------------//////****
  2   0238:59.55                                                       I
  1   0239:19.55      0     0     1     4                        <TL>  I-----------------------////***
  1   0239:39.55      0     0     0     1     0     3            <TL>  I-------------------------------/////////**
  1   0239:59.55                                                       I
  1   0240:19.55                                                       I
  2   0240:39.55      0     1     0     3                        <TM>  I-------------------/////////**
  2   0240:59.55                                                       I
  2   0241:19.55      0     1     0     3     1                  <TM>  I-------------------//////////*****
  2   0241:39.55      0     0     1     0     1                  <TM>  I-----------------------/////////****
  2   0241:59.55      0     0     0     0     0     2            <TL>  I---------------------------------------2
  2   0242:19.55                                                       I
  3   0242:39.55      0     0     2     3     1     1            <TM>  I---------------------/////////////********
  3   0242:59.55      0     0     0     0     0     2            <TN>  I33333333333333333333333333333333333333333
  3   0243:19.55                                                       I
  4   0243:39.55      0     0     1     3                        <TD>  I------------------------///***
  4   0243:59.55      0     0     1     0     1     2            <TL>  I---------------------////////////////****
  4   0244:04.13      0     0     0     1                        <TN>  I33333333333333333333333333333


  CURVE        MINIMUM         MAXIMUM
  MIN        1.7286E 00      2.6454E 01
  AVG        3.7757E 00      2.6636E 01
  MAX        5.4922E 00      3.0057E 01
```

Figure 14-26. User Swaps (Duration) Report

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982
                    DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0236:21.243 FRI 83-02-11.  INITIAL GMC TAPE #0276


                                                                  USER I/OS (DURATION)                                                PLOT-1

                                                                    -= MIN    /= AVG    *= MAX

                          DURATION IN INTERVAL (HS)                         I/O TIMES (HUNDREDTHS SEC)
USERS   TIME             <1   <2   <4   <8  <16  <32  >32  LINE 1.0E-01          1.0E 00          1.0E 01          1.0E 02
------------------------------------------------------------------Y----------------Y----------------Y----------------Y----------------Y
  0  0237:19.55  DATA NOT COLLECTED IN THIS INTERVAL.             I
  2  0237:39.55  76    0    4                           <TL>  I--------------//*********
  2  0237:59.55  66    0    5                           <TL>  I--------------///********
  2  0238:19.55  64    1    4                           <TL>  I--------------///********
  2  0238:39.55                                               I
  2  0238:59.55                                               I
  1  0239:19.55   1    0    1    2                      <TL>  I---------------////////////***
  1  0239:39.55   0    1    1    0    0    1            <TL>  I-------------------////////////////*******
  1  0239:59.55                                               I
  1  0240:19.55                                               I
  2  0240:39.55   0    1    0    1                      <TM>  I--------------------////***
  2  0240:59.55                                               I
  2  0241:19.55   0    1    2    2    1                 <TM>  I------------------////////////*****
  2  0241:39.55   0    1    0    3    1                 <TM>  I------------------////////////*****
  2  0241:59.55                                               I
  2  0242:19.55                                               I
  3  0242:39.55   0    3    1    2    0    1            <TM>  I----------------//////////////***********
  3  0242:59.55                                               I
  3  0243:19.55                                               I
  4  0243:39.55   0    1    1                           <TO>  I--------------------/*
  4  0243:59.55   1    1    1    1    1                 <TO>  I---------------///////////******
  4  0244:04.13                                               I


   CURVE        MINIMUM        MAXIMUM
   MIN        7.2160E-01     1.8677E 00
   AVG        1.0213E 00     9.4907E 00
   MAX        2.5036E 00     2.3888E 01
```

Figure 14-27. User - I/Os - Duration Report

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPT 1982

                DISTRIBUTION COLLECTED ON SYSTEM BSCC1  AT 0236:21.243 FRI 83-02-11.  INITIAL GMC TAPE #0276


                                        IN SMC WAIT (DURATION)                                        PLOT-15

                                          -= MIN    /= AVG    += MAX

                  WAITS IN INTERVAL (HS)                      WAIT TIMES (HUNDREDTHS SEC)
USERS   TIME        <1   <2   <4   <8  <16  <32  >32  LINE 1.0E 00     1.0E 01      1.0E 02      1.0E 03      1.0E 04
---------------------------------------------------------Y------------Y------------Y------------Y------------Y------------Y
  0  0237:19.55  DATA NOT COLLECTED IN THIS INTERVAL.     X
  2  0237:39.55                                           I
  2  0237:59.55                                           I
  2  0238:19.55                                           I
  2  0238:39.55                                           I
  2  0238:59.55                                           I
  1  0239:19.55                                           I
  1  0239:39.55                                           I
  1  0239:59.55                                           I
  1  0240:19.55                                           I
  2  0240:39.55                                           X
  2  0240:59.55                                           I
  2  0241:19.55                                           I
  2  0241:39.55                                           I
  2  0241:59.55                                           I
  2  0242:19.55                                           I
  3  0242:39.55                                           I
  3  0242:59.55                                           I
  3  0243:19.55                                           I
  4  0243:39.55                                           I
  4  0243:59.55                                           X
  4  0244:04.13                                           I


CURVE        MINIMUM         MAXIMUM
MIN        9.9000E 33     -9.9000E 33
AVG        9.9000E 33     -9.9000E 33
MAX        9.9000E 33     -9.9000E 33
```

Figure 14-28. In SMC WAIT (Duration) Report different data tapes should reveal similar problems in the same types of reports.

The next step in the resolution of a response problem will be a decision as whether the problem requires action at all. That is to say, are there any steps which can be taken to affect the delay which has been discovered? For example, if a user is being (or was) delayed due to an unusually high workload, there is really nothing which can be done to improve the situation without affecting processing during normal periods. Another example of this type of situation can be encountered by examining some of the TEARS reports which seem to indicate problems which are not problems at all. Certainly, the tunable parameters associated with TSS could be used to affect the balance between Fault queue wait time and Ready queue service time, but this 'solution' might cause more problems than it solves.

Before beginning the modification of the TSS and GCOS software to change the values reported by the TEARS system (and, hopefully, the response time experienced by TSS users), site personnel should first consider the cost of the loss represented by the response problem incurred. If it is not encountered except during unusual periods, the time and effort (in both human and machine resources) required to prevent it may not be justified. The process of analysing and preventing or minimizing response time problems is an expensive one, and should not be lightly undertaken.

The process involves the modification of the tunable parameters described in section 14.6.7.3 of this report. The degree of modification required, and the exact parameters which will need to be modified can be discovered only through experimentation. This experimentation must, of course, be carried out in the same environment as that under which the problem was discovered. It is unlikely that a single setting change will resolve a response problem. More likely, a number of settings will need to be tried for several parameters before an optimal setting discovered.

The next section of this report describes, in general terms, the process of this modification, and the types of parameters which should be modified to affect different problems.

14.6.7.5 Analysis of Response Time Problems. The resolution of response time problems through the modification of the tunable parameters described in section 14.6.7.3 of this report is possible if the problems are related to the TSS system itself. If the problems are due to user actions or some circumstance, such as disk space fragmentation, modification of these parameters will have little effect, and other types of action are indicated.

The sections within this final chapter are divided according to the subject areas covered in section 14.6.7.3. If modification of the parameters described in section 14.6.7.3 do not affect the problem encountered, it is probable that some user or system problem exists. Successful completion of these procedures will not be possible if the reader is not cognizant of the effect these modifications are having. Therefore, all modifications performed should be both preceded and followed by executions of the GMF/TSSM and TEARS. As more modifications are performed, and more trace data collected, the nature of the problem can be isolated to a great degree. For example, if lengthening the dispatches granted to the TSS executive does not affect response time, the problem will not involve the Sub-dispatch Fault queue. It will, at that point, not be helpful to increase the frequency of priority 'B' dispatches granted to TSS.

14.6.7.5.1 Priority 'B' Processing Problems. The TSS executive relies on .MDISP to grant it enough processor attention to completely service the entries in its Fault queue in a timely manner. If the executive is not getting sufficient attention, the response times for users performing allocation processes and I/O operations will be increased. Also, the report displaying the wait time for Fault queue service will show high values.

To resolve these problems, the number of times that TSS is able to service its Fault queue must be increased. First, increase the length of a TSS executive dispatch by modifying the patches to .MDISP and TSSA shown in section 14.6.7.3.1. This should lower the wait time for Fault queue service, as well as the total average response times experienced. Note also the slope of the average response times for users with core requests, which should also be lowered.

If this modification does affect the response time, but not to the degree desired, the user may wish to increase the number of times that TSS will be granted a priority 'B' dispatch, by lowering the value of bits 30-32 of location 1 in .MDISP.

14.6.7.5.2 TSS Executive File Problems. If the modification of TSS dispatch lengths does not affect the response time for users with core requests, the fault may lie with the TSS Swap files. Verify that there is a high degree of swapping being performed by examining the TSS subtraces frequency (section 14.6.7.4.2.9). Next, verify that the recommendations contained in section 14.6.7.3.2 have been carried out, and that the swap files are not contending with one another by being resident on the same device. Before modifying any of the TSS parameters associated with memory allocation or swap space management, first implement these recommendations and recollect trace data.

If the implementation of the suggested placement for TSS files does not affect the reports described above, the problem may lie in the size of the files. Examine the trace frequency report (section 14.6.7.4.2.9) for trace type 2, which will be produced if the TSS swap files are too small. If encountered, increase the size for each of the active swap files and recollect trace data. If not encountered, examine the counts for the traces associated with urgent user processing. These traces will indicate a problem with subsystem memory management rather that the swap files themselves. Finally, ensure that the TSS program files are not on the same device by examining the system boot deck.

14.6.7.5.3 TSS Swap File Processing Problems. The symptoms of problems with the algorithms which control the swapping of user subsystems are: numerous swaps, force swapping, urgent user allocations, as well as the occurrence of those GMF/TSSM traces which record TSS size changes. These problems are caused by contention for a limited amount of core storage. The user should first examine the size limitations placed on TSS before performing any modifications to the swapping algorithms.

The amount of core storage which can be given to TSS is, of course, limited by the amount of core available to the system as well as that required by other processes. If the loading of the system under examination precludes resolution of swap file processing problems through expansion of the executive core space, the analyst may begin the optimization process by modifying the definitions of large subsystems through the directions shown in section 14.6.7.3.7.2 This modification will limit the number of times that force swaps occur due to urgent allocation processes. This should be followed by a recollection of trace data, as problems with the files themselves is much less common than high contention among user subsystems for limited core space.

Examine the number of occurrences of the exception message (section 14.6.7.4.3.1) SWAP SPACE DENIED. Any occurrence of this message should cause modifications to the minimum sizes, and amount to grow, for these swap files (section 14.6.7.3.3). Increase these values and recollect trace information until this exception fails to occur.

14.6.7.5.4 Subsystem Accounting Problems. To evaluate the effect of subsystem accounting on the response times being experienced, disable this feature as shown in section 14.6.7.3.4, and recollect the trace data. Examine each of the reports for any difference. Little will be noted. Unless some GCOS system problem is being encountered, this processing will have little effect. Unless the data collected by subsystem accounting is required by the site, however, the overhead represented by this processing should be turned off.

14.6.7.5.5 UST Memory Management Problems. The only problem which is likely to be encountered related to UST management involves the delay imposed before the release of abnormally disconnected UST entries.

Examine the subtraces encountered report (section 14.6.7.4.2.9) for the following trace types:

- 22 - Place User in Reconnect Mode (Abnormal disconnect has just occurred)
- 71 - UST area increase by 1K
- 97 - UST compressions
- 99 - UST area decrease by 1K

The count for trace type 71 will give an indication of how often these abnormal disconnects occur. Trace types 71, 97, and 99 will give an indication of how dynamic the load faced by TSS is during a tracing period. Examine the event log and the response times reports to gain an understanding of the number of times the number of active users approaches the maximum (section 14.6.7.3.5). Finally, look for the number of reconnect USTs which are eventually recovered. The reconnect will cause the generation of trace type 58.

Typically, the number of USTs which are placed in reconnect mode is much larger than the number which are eventually reconnected to. If the system is extremely dynamic (traces 71, 97, 99), or if the number of active users approaches the maximum, the timer for these reconnect USTs should be lowered, causing them to become available sooner. The damper on DRUN sessions can also be used to lower the frequency of these trace types at sites which employ this deferred command file processing facility.

14.6.7.5.6 Subsystem Memory Management Problems

Due to the limited nature of core storage which is available to TSS, and the large size of the subsystems which are usually requested by TSS users, the contention for space in the subsystem memory area is usually high. To manage this space, TSS uses algorithms which tend to discriminate against large memory allocation requests. The degree to which TSS will discriminate against these systems is adjustable, and will provide the best mechanism to lower the response times experienced by users who require memory allocation (section 14.6.7.4.2.4). The first parameter which should be changed is the size definition for a large subsystem. The optimal setting for this parameter is twice the size of average subsystem size (computed based on usage). Unfortunately, this is not reported by the TEARS system. Experimentation will reveal that 20K is a good value for this cell.

If modification of the definition of a large subsystem does not yield the desired results against the wait times experienced by users requiring memory allocation, the next parameter to modify will be the wait time multiplier, which will cause a large memory allocation to wait longer before causing swaps to occur.

Ironically, the suggested modifications to decrease the wait time for memory allocations will cause increases in certain allocation times. The effect is, however, positive. By discouraging large subsystem allocations, smaller subsystems are able to load faster, and terminate, leaving space open to complete the large allocation. This principle has been employed on heavily loaded systems with very positive results.

The other factor which will affect subsystem memory allocations, and hence, response time, involves the growth of the TSS executive within the bounds established by site options. The amount of core which will be requested by the executive should be adjusted based on the occurrence of trace types 72 and 74 (GEMORE memory, release memory). While performing a GEMORE, TSS is vulnerable, and is unable to service its Fault queue, causing delays for users who do not require memory.

By increasing the amount which will be requested on each GEMORE, the number of these size increases can be minimized. When performing this type of modification, corresponding changes should be made in the size of a memory reduction. The amount of core to release should be smaller, but an even divisor of the amount of core which is attached. Thus, TSS will quickly get more core, but will give it back slower. These patches are described in section 14.6.7.3.7.

14.6.7.5.7 TSS Executive Processing Problems. Trace type 3s (section 14.6.7.4.2.9) should cause modifications to the number of SSAs reserved by TSS during its initialization processing, according to the patches described in section 14.6.7.3.7.